## All of

# Nonparametric 

 Statistics
## Springer Texts in Statistics

Advisors:
George Casella Stephen Fienberg Ingram Olkin

## Springer Texts in Statistics

Alfred: Elements of Statistics for the Life and Social Sciences
Berger: An Introduction to Probability and Stochastic Processes
Bilodeau and Brenner: Theory of Multivariate Statistics
Blom: Probability and Statistics: Theory and Applications
Brockwell and Davis: Introduction to Times Series and Forecasting, Second Edition
Chow and Teicher: Probability Theory: Independence, Interchangeability, Martingales, Third Edition
Christensen: Advanced Linear Modeling: Multivariate, Time Series, and Spatial Data-Nonparametric Regression and Response Surface Maximization, Second Edition
Christensen: Log-Linear Models and Logistic Regression, Second Edition
Christensen: Plane Answers to Complex Questions: The Theory of Linear Models, Third Edition
Creighton: A First Course in Probability Models and Statistical Inference
Davis: Statistical Methods for the Analysis of Repeated Measurements
Dean and Voss: Design and Analysis of Experiments
du Toit, Steyn, and Stumpf: Graphical Exploratory Data Analysis
Durrett: Essentials of Stochastic Processes
Edwards: Introduction to Graphical Modelling, Second Edition
Finkelstein and Levin: Statistics for Lawyers
Flury: A First Course in Multivariate Statistics
Jobson: Applied Multivariate Data Analysis, Volume I: Regression and Experimental Design
Jobson: Applied Multivariate Data Analysis, Volume II: Categorical and Multivariate Methods
Kalbfleisch: Probability and Statistical Inference, Volume I: Probability, Second Edition
Kalbfleisch: Probability and Statistical Inference, Volume II: Statistical Inference, Second Edition
Karr: Probability
Keyfitz: Applied Mathematical Demography, Second Edition
Kiefer: Introduction to Statistical Inference
Kokoska and Nevison: Statistical Tables and Formulae
Kulkarni: Modeling, Analysis, Design, and Control of Stochastic Systems
Lange: Applied Probability
Lehmann: Elements of Large-Sample Theory
Lehmann: Testing Statistical Hypotheses, Second Edition
Lehmann and Casella: Theory of Point Estimation, Second Edition
Lindman: Analysis of Variance in Experimental Design
Lindsey: Applying Generalized Linear Models

Larry Wasserman

## All of Nonparametric Statistics

With 52 Illustrations

Larry Wasserman<br>Department of Statistics<br>Carnegie Mellon University<br>Pittsburgh, PA 15213-3890<br>USA<br>larry@stat.cmu.edu

## Editorial Board

George Casella
Department of Statistics
University of Florida
Gainesville, FL 32611-8545
USA

Stephen Fienberg<br>Department of Statistics<br>Carnegie Mellon University<br>Pittsburgh, PA 15213-3890<br>USA

Ingram Olkin<br>Department of Statistics<br>Stanford University<br>Stanford, CA 94305<br>USA

Library of Congress Control Number: 2005925603
ISBN-10: 0-387-25145-6
ISBN-13: 978-0387-25145-5
Printed on acid-free paper.
© 2006 Springer Science+Business Media, Inc.
All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.
The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America. (MVY)

## 987654321

springeronline.com

To Isa

## Preface

There are many books on various aspects of nonparametric inference such as density estimation, nonparametric regression, bootstrapping, and wavelets methods. But it is hard to find all these topics covered in one place. The goal of this text is to provide readers with a single book where they can find a brief account of many of the modern topics in nonparametric inference.

The book is aimed at master's-level or Ph.D.-level statistics and computer science students. It is also suitable for researchers in statistics, machine learning and data mining who want to get up to speed quickly on modern nonparametric methods. My goal is to quickly acquaint the reader with the basic concepts in many areas rather than tackling any one topic in great detail. In the interest of covering a wide range of topics, while keeping the book short, I have opted to omit most proofs. Bibliographic remarks point the reader to references that contain further details. Of course, I have had to choose topics to include and to omit, the title notwithstanding. For the most part, I decided to omit topics that are too big to cover in one chapter. For example, I do not cover classification or nonparametric Bayesian inference.

The book developed from my lecture notes for a half-semester (20 hours) course populated mainly by master's-level students. For Ph.D.-level students, the instructor may want to cover some of the material in more depth and require the students to fill in proofs of some of the theorems. Throughout, I have attempted to follow one basic principle: never give an estimator without giving a confidence set.

The book has a mixture of methods and theory. The material is meant to complement more method-oriented texts such as Hastie et al. (2001) and Ruppert et al. (2003).

After the Introduction in Chapter 1, Chapters 2 and 3 cover topics related to the empirical CDF such as the nonparametric delta method and the bootstrap. Chapters 4 to 6 cover basic smoothing methods. Chapters 7 to 9 have a higher theoretical content and are more demanding. The theory in Chapter 7 lays the foundation for the orthogonal function methods in Chapters 8 and 9. Chapter 10 surveys some of the omitted topics.

I assume that the reader has had a course in mathematical statistics such as Casella and Berger (2002) or Wasserman (2004). In particular, I assume that the following concepts are familiar to the reader: distribution functions, convergence in probability, convergence in distribution, almost sure convergence, likelihood functions, maximum likelihood, confidence intervals, the delta method, bias, mean squared error, and Bayes estimators. These background concepts are reviewed briefly in Chapter 1.

Data sets and code can be found at:

```
www.stat.cmu.edu/~larry/all-of-nonpar
```

I need to make some disclaimers. First, the topics in this book fall under the rubric of "modern nonparametrics." The omission of traditional methods such as rank tests and so on is not intended to belittle their importance. Second, I make heavy use of large-sample methods. This is partly because I think that statistics is, largely, most successful and useful in large-sample situations, and partly because it is often easier to construct large-sample, nonparametric methods. The reader should be aware that large-sample methods can, of course, go awry when used without appropriate caution.

I would like to thank the following people for providing feedback and suggestions: Larry Brown, Ed George, John Lafferty, Feng Liang, Catherine Loader, Jiayang Sun, and Rob Tibshirani. Special thanks to some readers who provided very detailed comments: Taeryon Choi, Nils Hjort, Woncheol Jang, Chris Jones, Javier Rojo, David Scott, and one anonymous reader. Thanks also go to my colleague Chris Genovese for lots of advice and for writing the $\mathrm{LA}_{\mathrm{E}} \mathrm{X}$ macros for the layout of the book. I am indebted to John Kimmel, who has been supportive and helpful and did not rebel against the crazy title. Finally, thanks to my wife Isabella Verdinelli for suggestions that improved the book and for her love and support.

Larry Wasserman
Pittsburgh, Pennsylvania
July 2005

## Contents

1 Introduction ..... 1
1.1 What Is Nonparametric Inference? ..... 1
1.2 Notation and Background ..... 2
1.3 Confidence Sets ..... 5
1.4 Useful Inequalities ..... 8
1.5 Bibliographic Remarks ..... 10
1.6 Exercises ..... 10
2 Estimating the CDF and Statistical Functionals ..... 13
2.1 The CDF ..... 13
2.2 Estimating Statistical Functionals ..... 15
2.3 Influence Functions ..... 18
2.4 Empirical Probability Distributions ..... 21
2.5 Bibliographic Remarks ..... 23
2.6 Appendix ..... 23
2.7 Exercises ..... 24
3 The Bootstrap and the Jackknife ..... 27
3.1 The Jackknife ..... 27
3.2 The Bootstrap ..... 30
3.3 Parametric Bootstrap ..... 31
3.4 Bootstrap Confidence Intervals ..... 32
3.5 Some Theory ..... 35
3.6 Bibliographic Remarks ..... 37
3.7 Appendix ..... 37
3.8 Exercises ..... 39
4 Smoothing: General Concepts ..... 43
4.1 The Bias-Variance Tradeoff ..... 50
4.2 Kernels ..... 55
4.3 Which Loss Function? ..... 57
4.4 Confidence Sets ..... 57
4.5 The Curse of Dimensionality ..... 58
4.6 Bibliographic Remarks ..... 59
4.7 Exercises ..... 59
5 Nonparametric Regression ..... 61
5.1 Review of Linear and Logistic Regression ..... 63
5.2 Linear Smoothers ..... 66
5.3 Choosing the Smoothing Parameter ..... 68
5.4 Local Regression ..... 71
5.5 Penalized Regression, Regularization and Splines ..... 81
5.6 Variance Estimation ..... 85
5.7 Confidence Bands ..... 89
5.8 Average Coverage ..... 94
5.9 Summary of Linear Smoothing ..... 95
5.10 Local Likelihood and Exponential Families ..... 96
5.11 Scale-Space Smoothing ..... 99
5.12 Multiple Regression ..... 100
5.13 Other Issues ..... 111
5.14 Bibliographic Remarks ..... 119
5.15 Appendix ..... 119
5.16 Exercises ..... 120
6 Density Estimation ..... 125
6.1 Cross-Validation ..... 126
6.2 Histograms ..... 127
6.3 Kernel Density Estimation ..... 131
6.4 Local Polynomials ..... 137
6.5 Multivariate Problems ..... 138
6.6 Converting Density Estimation Into Regression ..... 139
6.7 Bibliographic Remarks ..... 140
6.8 Appendix ..... 140
6.9 Exercises ..... 142
7 Normal Means and Minimax Theory ..... 145
7.1 The Normal Means Model ..... 145
7.2 Function Spaces ..... 147
7.3 Connection to Regression and Density Estimation ..... 149
7.4 Stein's Unbiased Risk Estimator (Sure) ..... 150
7.5 Minimax Risk and Pinsker's Theorem ..... 153
7.6 Linear Shrinkage and the James-Stein Estimator ..... 155
7.7 Adaptive Estimation Over Sobolev Spaces ..... 158
7.8 Confidence Sets ..... 159
7.9 Optimality of Confidence Sets ..... 166
7.10 Random Radius Bands? ..... 170
7.11 Penalization, Oracles and Sparsity ..... 171
7.12 Bibliographic Remarks ..... 172
7.13 Appendix ..... 173
7.14 Exercises ..... 180
8 Nonparametric Inference Using Orthogonal Functions ..... 183
8.1 Introduction ..... 183
8.2 Nonparametric Regression ..... 183
8.3 Irregular Designs ..... 190
8.4 Density Estimation ..... 192
8.5 Comparison of Methods ..... 193
8.6 Tensor Product Models ..... 193
8.7 Bibliographic Remarks ..... 194
8.8 Exercises ..... 194
9 Wavelets and Other Adaptive Methods ..... 197
9.1 Haar Wavelets ..... 199
9.2 Constructing Wavelets ..... 203
9.3 Wavelet Regression ..... 206
9.4 Wavelet Thresholding ..... 208
9.5 Besov Spaces ..... 211
9.6 Confidence Sets ..... 214
9.7 Boundary Corrections and Unequally Spaced Data ..... 215
9.8 Overcomplete Dictionaries ..... 215
9.9 Other Adaptive Methods ..... 216
9.10 Do Adaptive Methods Work? ..... 220
9.11 Bibliographic Remarks ..... 221
9.12 Appendix ..... 221
9.13 Exercises ..... 223
10 Other Topics ..... 227
10.1 Measurement Error ..... 227
10.2 Inverse Problems ..... 233
10.3 Nonparametric Bayes ..... 235
10.4 Semiparametric Inference ..... 235
10.5 Correlated Errors ..... 236
10.6 Classification ..... 236
10.7 Sieves ..... 237
10.8 Shape-Restricted Inference ..... 237
10.9 Testing ..... 238
10.10Computational Issues ..... 240
10.11Exercises ..... 240
Bibliography ..... 243
List of Symbols ..... 259
Table of Distributions ..... 261
Index ..... 263

## 1

## Introduction

In this chapter we briefly describe the types of problems with which we will be concerned. Then we define some notation and review some basic concepts from probability theory and statistical inference.

### 1.1 What Is Nonparametric Inference?

The basic idea of nonparametric inference is to use data to infer an unknown quantity while making as few assumptions as possible. Usually, this means using statistical models that are infinite-dimensional. Indeed, a better name for nonparametric inference might be infinite-dimensional inference. But it is difficult to give a precise definition of nonparametric inference, and if I did venture to give one, no doubt I would be barraged with dissenting opinions.

For the purposes of this book, we will use the phrase nonparametric inference to refer to a set of modern statistical methods that aim to keep the number of underlying assumptions as weak as possible. Specifically, we will consider the following problems:

1. (Estimating the distribution function). Given an IID sample $X_{1}, \ldots, X_{n} \sim$ $F$, estimate the CDF $F(x)=\mathbb{P}(X \leq x)$. (Chapter 2.)
2. (Estimating functionals). Given an IID sample $X_{1}, \ldots, X_{n} \sim F$, estimate a functional $T(F)$ such as the mean $T(F)=\int x d F(x)$. (Chapters 2 and 3.)
3. (Density estimation). Given an IID sample $X_{1}, \ldots, X_{n} \sim F$, estimate the density $f(x)=F^{\prime}(x)$. (Chapters 4,6 and 8 .)
4. (Nonparametric regression or curve estimation). Given $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ estimate the regression function $r(x)=\mathbb{E}(Y \mid X=x)$. (Chapters 4, 5, 8 and 9.)
5. (Normal means). Given $Y_{i} \sim N\left(\theta_{i}, \sigma^{2}\right), i=1, \ldots, n$, estimate $\theta=$ $\left(\theta_{1}, \ldots, \theta_{n}\right)$. This apparently simple problem turns out to be very complex and provides a unifying basis for much of nonparametric inference. (Chapter 7.)

In addition, we will discuss some unifying theoretical principles in Chapter 7. We consider a few miscellaneous problems in Chapter 10, such as measurement error, inverse problems and testing.

Typically, we will assume that distribution $F$ (or density $f$ or regression function $r$ ) lies in some large set $\mathfrak{F}$ called a statistical model. For example, when estimating a density $f$, we might assume that

$$
f \in \mathfrak{F}=\left\{g: \quad \int\left(g^{\prime \prime}(x)\right)^{2} d x \leq c^{2}\right\}
$$

which is the set of densities that are not "too wiggly."

### 1.2 Notation and Background

Here is a summary of some useful notation and background. See also Table 1.1.

Let $a(x)$ be a function of $x$ and let $F$ be a cumulative distribution function. If $F$ is absolutely continuous, let $f$ denote its density. If $F$ is discrete, let $f$ denote instead its probability mass function. The mean of $a$ is

$$
\mathbb{E}(a(X))=\int a(x) d F(x) \equiv \begin{cases}\int a(x) f(x) d x & \text { continuous case } \\ \sum_{j} a\left(x_{j}\right) f\left(x_{j}\right) & \text { discrete case. }\end{cases}
$$

Let $\mathbb{V}(X)=\mathbb{E}(X-\mathbb{E}(X))^{2}$ denote the variance of a random variable. If $X_{1}, \ldots, X_{n}$ are $n$ observations, then $\int a(x) d \widehat{F}_{n}(x)=n^{-1} \sum_{i} a\left(X_{i}\right)$ where $\widehat{F}_{n}$ is the empirical distribution that puts mass $1 / n$ at each observation $X_{i}$.

| $\frac{\text { Symbol }}{x_{n}=o\left(a_{n}\right)}$ | $\underline{\text { Definition }}$ |
| :--- | :--- |
| $x_{n}=O\left(a_{n}\right)$ | $\left\|x_{n} / a_{n}\right\|$ is bounded for all large $n$ |
| $a_{n} \sim b_{n}$ | $a_{n} / b_{n} \rightarrow 1$ as $n \rightarrow \infty$ |
| $a_{n} \asymp b_{n}$ | $a_{n} / b_{n}$ and $b_{n} / a_{n}$ are bounded for all large $n$ |
| $X_{n} \rightsquigarrow X$ | convergence in distribution |
| $X_{n} \xrightarrow{\mathrm{P}} X$ | convergence in probability |
| $X_{n} \xrightarrow{\text { a.s. }} X$ | almost sure convergence |
| $\widehat{\theta}_{n}$ | estimator of parameter $\theta$ |
| bias | $\mathbb{E}\left(\widehat{\theta}_{n}\right)-\theta$ |
| se | $\sqrt{\mathbb{V}\left(\widehat{\theta}_{n}\right)}$ (standard error) |
| $\widehat{\text { se }}$ | estimated standard error |
| MSE | $\mathbb{E}\left(\widehat{\theta}_{n}-\theta\right)^{2}$ (mean squared error) |
| $\Phi$ | CDF of a standard Normal random variable |
| $z_{\alpha}$ | $\Phi^{-1}(1-\alpha)$ |
|  |  |
|  | TABLE 1.1. Some useful notation. |
|  |  |

Brief Review of Probability. The sample space $\Omega$ is the set of possible outcomes of an experiment. Subsets of $\Omega$ are called events. A class of events $\mathcal{A}$ is called a $\sigma$-field if (i) $\emptyset \in \mathcal{A}$, (ii) $A \in \mathcal{A}$ implies that $A^{c} \in \mathcal{A}$ and (iii) $A_{1}, A_{2}, \ldots, \in \mathcal{A}$ implies that $\bigcup_{i=1}^{\infty} A_{i} \in \mathcal{A}$. A probability measure is a function $\mathbb{P}$ defined on a $\sigma$-field $\mathcal{A}$ such that $\mathbb{P}(A) \geq 0$ for all $A \in \mathcal{A}, \mathbb{P}(\Omega)=1$ and if $A_{1}, A_{2}, \ldots \in \mathcal{A}$ are disjoint then

$$
\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_{i}\right)=\sum_{i=1}^{\infty} \mathbb{P}\left(A_{i}\right)
$$

The triple $(\Omega, \mathcal{A}, \mathbb{P})$ is called a probability space. A random variable is a map $X: \Omega \rightarrow \mathbb{R}$ such that, for every real $x,\{\omega \in \Omega: X(\omega) \leq x\} \in \mathcal{A}$.

A sequence of random variables $X_{n}$ converges in distribution (or converges weakly) to a random variable $X$, written $X_{n} \rightsquigarrow X$, if

$$
\begin{equation*}
\mathbb{P}\left(X_{n} \leq x\right) \rightarrow \mathbb{P}(X \leq x) \tag{1.1}
\end{equation*}
$$

as $n \rightarrow \infty$, at all points $x$ at which the CDF

$$
\begin{equation*}
F(x)=\mathbb{P}(X \leq x) \tag{1.2}
\end{equation*}
$$

is continuous. A sequence of random variables $X_{n}$ converges in probability to a random variable $X$, written $X_{n} \xrightarrow{\mathrm{P}} X$, if,

$$
\begin{equation*}
\text { for every } \epsilon>0, \quad \mathbb{P}\left(\left|X_{n}-X\right|>\epsilon\right) \rightarrow 0 \quad \text { as } n \rightarrow \infty . \tag{1.3}
\end{equation*}
$$

A sequence of random variables $X_{n}$ converges almost surely to a random variable $X$, written $X_{n} \xrightarrow{\text { a.s. }} X$, if

$$
\begin{equation*}
\mathbb{P}\left(\lim _{n \rightarrow \infty}\left|X_{n}-X\right|=0\right)=1 \tag{1.4}
\end{equation*}
$$

The following implications hold:

$$
\begin{equation*}
X_{n} \xrightarrow{\text { a.s. }} X \quad \text { implies that } \quad X_{n} \xrightarrow{\mathrm{P}} X \quad \text { implies that } \quad X_{n} \rightsquigarrow X . \tag{1.5}
\end{equation*}
$$

Let $g$ be a continuous function. Then, according to the continuous mapping theorem,

$$
\begin{array}{lll}
X_{n} \rightsquigarrow X & \text { implies that } & g\left(X_{n}\right) \rightsquigarrow g(X) \\
X_{n} \xrightarrow{\mathrm{P}} X & \text { implies that } & g\left(X_{n}\right) \xrightarrow{\mathrm{P}} g(X) \\
X_{n} \xrightarrow{\text { a.s. }} X & \text { implies that } & g\left(X_{n}\right) \xrightarrow{\text { a.s. }} g(X)
\end{array}
$$

According to Slutsky's theorem, if $X_{n} \rightsquigarrow X$ and $Y_{n} \rightsquigarrow c$ for some constant $c$, then $X_{n}+Y_{n} \rightsquigarrow X+c$ and $X_{n} Y_{n} \rightsquigarrow c X$.

Let $X_{1}, \ldots, X_{n} \sim F$ be IID. The weak law of large numbers says that if $\mathbb{E}\left|g\left(X_{1}\right)\right|<\infty$, then $n^{-1} \sum_{i=1}^{n} g\left(X_{i}\right) \xrightarrow{\mathrm{P}} \mathbb{E}\left(g\left(X_{1}\right)\right)$. The strong law of large numbers says that if $\mathbb{E}\left|g\left(X_{1}\right)\right|<\infty$, then $n^{-1} \sum_{i=1}^{n} g\left(X_{i}\right) \xrightarrow{\text { a.s. }} \mathbb{E}\left(g\left(X_{1}\right)\right)$.

The random variable $Z$ has a standard Normal distribution if it has density $\phi(z)=(2 \pi)^{-1 / 2} e^{-z^{2} / 2}$ and we write $Z \sim N(0,1)$. The CDF is denoted by $\Phi(z)$. The $\alpha$ upper quantile is denoted by $z_{\alpha}$. Thus, if $Z \sim N(0,1)$, then $\mathbb{P}\left(Z>z_{\alpha}\right)=\alpha$.

If $\mathbb{E}\left(g^{2}\left(X_{1}\right)\right)<\infty$, the central limit theorem says that

$$
\begin{equation*}
\sqrt{n}\left(\bar{Y}_{n}-\mu\right) \rightsquigarrow N\left(0, \sigma^{2}\right) \tag{1.6}
\end{equation*}
$$

where $Y_{i}=g\left(X_{i}\right), \mu=\mathbb{E}\left(Y_{1}\right), \bar{Y}_{n}=n^{-1} \sum_{i=1}^{n} Y_{i}$ and $\sigma^{2}=\mathbb{V}\left(Y_{1}\right)$. In general, if

$$
\frac{\left(X_{n}-\mu\right)}{\widehat{\sigma}_{n}} \rightsquigarrow N(0,1)
$$

then we will write

$$
\begin{equation*}
X_{n} \approx N\left(\mu, \widehat{\sigma}_{n}^{2}\right) \tag{1.7}
\end{equation*}
$$

According to the delta method, if $g$ is differentiable at $\mu$ and $g^{\prime}(\mu) \neq 0$ then

$$
\begin{equation*}
\sqrt{n}\left(X_{n}-\mu\right) \rightsquigarrow N\left(0, \sigma^{2}\right) \Longrightarrow \sqrt{n}\left(g\left(X_{n}\right)-g(\mu)\right) \rightsquigarrow N\left(0,\left(g^{\prime}(\mu)\right)^{2} \sigma^{2}\right) . \tag{1.8}
\end{equation*}
$$

A similar result holds in the vector case. Suppose that $X_{n}$ is a sequence of random vectors such that $\sqrt{n}\left(X_{n}-\mu\right) \rightsquigarrow N(0, \Sigma)$, a multivariate, mean 0
normal with covariance matrix $\Sigma$. Let $g$ be differentiable with gradient $\nabla g$ such that $\nabla_{\mu} \neq 0$ where $\nabla_{\mu}$ is $\nabla g$ evaluated at $\mu$. Then

$$
\begin{equation*}
\sqrt{n}\left(g\left(X_{n}\right)-g(\mu)\right) \rightsquigarrow N\left(0, \nabla_{\mu}^{T} \Sigma \nabla_{\mu}\right) . \tag{1.9}
\end{equation*}
$$

Statistical Concepts. Let $\mathfrak{F}=\{f(x ; \theta): \theta \in \Theta\}$ be a parametric model satisfying appropriate regularity conditions. The likelihood function based on IID observations $X_{1}, \ldots, X_{n}$ is

$$
\mathcal{L}_{n}(\theta)=\prod_{i=1}^{n} f\left(X_{i} ; \theta\right)
$$

and the $\log$-likelihood function is $\ell_{n}(\theta)=\log \mathcal{L}_{n}(\theta)$. The maximum likelihood estimator, or MLE $\widehat{\theta}_{n}$, is the value of $\theta$ that maximizes the likelihood. The score function is $s(X ; \theta)=\partial \log f(x ; \theta) / \partial \theta$. Under appropriate regularity conditions, the score function satisfies $\mathbb{E}_{\theta}(s(X ; \theta))=\int s(x ; \theta) f(x ; \theta) d x=0$. Also,

$$
\sqrt{n}\left(\widehat{\theta}_{n}-\theta\right) \rightsquigarrow N\left(0, \tau^{2}(\theta)\right)
$$

where $\tau^{2}(\theta)=1 / I(\theta)$ and

$$
I(\theta)=\mathbb{V}_{\theta}(s(x ; \theta))=\mathbb{E}_{\theta}\left(s^{2}(x ; \theta)\right)=-\mathbb{E}_{\theta}\left(\frac{\partial^{2} \log f(x ; \theta)}{\partial \theta^{2}}\right)
$$

is the Fisher information. Also,

$$
\frac{\left(\widehat{\theta}_{n}-\theta\right)}{\widehat{\mathrm{se}}} \rightsquigarrow N(0,1)
$$

where $\widehat{\mathrm{se}}^{2}=1 /\left(n I\left(\widehat{\theta}_{n}\right)\right)$. The Fisher information $I_{n}$ from $n$ observations satisfies $I_{n}(\theta)=n I(\theta)$; hence we may also write $\widehat{\mathrm{se}}^{2}=1 /\left(I_{n}\left(\widehat{\theta}_{n}\right)\right)$.

The bias of an estimator $\widehat{\theta}_{n}$ is $\mathbb{E}(\widehat{\theta})-\theta$ and the the mean squared error MSE is MSE $=\mathbb{E}(\widehat{\theta}-\theta)^{2}$. The bias-variance decomposition for the MSE of an estimator $\widehat{\theta}_{n}$ is

$$
\begin{equation*}
\mathrm{MSE}=\operatorname{bias}^{2}\left(\widehat{\theta}_{n}\right)+\mathbb{V}\left(\widehat{\theta}_{n}\right) \tag{1.10}
\end{equation*}
$$

### 1.3 Confidence Sets

Much of nonparametric inference is devoted to finding an estimator $\widehat{\theta}_{n}$ of some quantity of interest $\theta$. Here, for example, $\theta$ could be a mean, a density or a regression function. But we also want to provide confidence sets for these quantities. There are different types of confidence sets, as we now explain.

Let $\mathfrak{F}$ be a class of distribution functions $F$ and let $\theta$ be some quantity of interest. Thus, $\theta$ might be $F$ itself, or $F^{\prime}$ or the mean of $F$, and so on. Let $C_{n}$ be a set of possible values of $\theta$ which depends on the data $X_{1}, \ldots, X_{n}$. To emphasize that probability statements depend on the underlying $F$ we will sometimes write $\mathbb{P}_{F}$.
1.11 Definition. $C_{n}$ is a finite sample $1-\alpha$ confidence set if

$$
\begin{equation*}
\inf _{F \in \mathfrak{F}} \mathbb{P}_{F}\left(\theta \in C_{n}\right) \geq 1-\alpha \quad \text { for all } n \tag{1.12}
\end{equation*}
$$

$C_{n}$ is a uniform asymptotic $1-\alpha$ confidence set if

$$
\begin{equation*}
\liminf _{n \rightarrow \infty} \inf _{F \in \mathfrak{F}} \mathbb{P}_{F}\left(\theta \in C_{n}\right) \geq 1-\alpha \tag{1.13}
\end{equation*}
$$

$C_{n}$ is a pointwise asymptotic $1-\alpha$ confidence set $i f$,

$$
\begin{equation*}
\text { for every } F \in \mathfrak{F}, \quad \liminf _{n \rightarrow \infty} \mathbb{P}_{F}\left(\theta \in C_{n}\right) \geq 1-\alpha \tag{1.14}
\end{equation*}
$$

If $\|\cdot\|$ denotes some norm and $\widehat{f}_{n}$ is an estimate of $f$, then a confidence ball for $f$ is a confidence set of the form

$$
\begin{equation*}
C_{n}=\left\{f \in \mathfrak{F}:\left\|f-\widehat{f}_{n}\right\| \leq s_{n}\right\} \tag{1.15}
\end{equation*}
$$

where $s_{n}$ may depend on the data. Suppose that $f$ is defined on a set $\mathcal{X}$. A pair of functions $(\ell, u)$ is a $1-\alpha$ confidence band or confidence envelope if

$$
\begin{equation*}
\inf _{f \in \mathfrak{F}} \mathbb{P}(\ell(x) \leq f(x) \leq u(x) \text { for all } x \in \mathcal{X}) \geq 1-\alpha \tag{1.16}
\end{equation*}
$$

Confidence balls and bands can be finite sample, pointwise asymptotic and uniform asymptotic as above. When estimating a real-valued quantity instead of a function, $C_{n}$ is just an interval and we call $C_{n}$ a confidence interval.

Ideally, we would like to find finite sample confidence sets. When this is not possible, we try to construct uniform asymptotic confidence sets. The last resort is a pointwise asymptotic confidence interval. If $C_{n}$ is a uniform asymptotic confidence set, then the following is true: for any $\delta>0$ there exists an $n(\delta)$ such that the coverage of $C_{n}$ is at least $1-\alpha-\delta$ for all $n>n(\delta)$. With a pointwise asymptotic confidence set, there may not exist a finite $n(\delta)$. In this case, the sample size at which the confidence set has coverage close to $1-\alpha$ will depend on $f$ (which we don't know).
1.17 Example. Let $X_{1}, \ldots, X_{n} \sim \operatorname{Bernoulli}(p)$. A pointwise asymptotic $1-\alpha$ confidence interval for $p$ is

$$
\begin{equation*}
\widehat{p}_{n} \pm z_{\alpha / 2} \sqrt{\frac{\widehat{p}_{n}\left(1-\widehat{p}_{n}\right)}{n}} \tag{1.18}
\end{equation*}
$$

where $\widehat{p}_{n}=n^{-1} \sum_{i=1}^{n} X_{i}$. It follows from Hoeffding's inequality (1.24) that a finite sample confidence interval is

$$
\begin{equation*}
\hat{p}_{n} \pm \sqrt{\frac{1}{2 n} \log \left(\frac{2}{\alpha}\right)} \tag{1.19}
\end{equation*}
$$

1.20 Example (Parametric models). Let

$$
\mathfrak{F}=\{f(x ; \theta): \theta \in \Theta\}
$$

be a parametric model with scalar parameter $\theta$ and let $\widehat{\theta}_{n}$ be the maximum likelihood estimator, the value of $\theta$ that maximizes the likelihood function

$$
\mathcal{L}_{n}(\theta)=\prod_{i=1}^{n} f\left(X_{i} ; \theta\right)
$$

Recall that under suitable regularity assumptions,

$$
\widehat{\theta}_{n} \approx N\left(\theta, \widehat{\mathrm{se}}^{2}\right)
$$

where

$$
\widehat{\mathrm{se}}=\left(I_{n}\left(\widehat{\theta}_{n}\right)\right)^{-1 / 2}
$$

is the estimated standard error of $\widehat{\theta}_{n}$ and $I_{n}(\theta)$ is the Fisher information. Then

$$
\widehat{\theta}_{n} \pm z_{\alpha / 2} \widehat{\operatorname{se}}
$$

is a pointwise asymptotic confidence interval. If $\tau=g(\theta)$ we can get an asymptotic confidence interval for $\tau$ using the delta method. The mLe for $\tau$ is $\widehat{\tau}_{n}=g\left(\widehat{\theta}_{n}\right)$. The estimated standard error for $\tau$ is $\widehat{\operatorname{se}}\left(\widehat{\tau}_{n}\right)=\widehat{\operatorname{se}}\left(\widehat{\theta}_{n}\right)\left|g^{\prime}\left(\widehat{\theta}_{n}\right)\right|$. The confidence interval for $\tau$ is

$$
\widehat{\tau}_{n} \pm z_{\alpha / 2} \widehat{\operatorname{se}}\left(\widehat{\tau}_{n}\right)=\widehat{\tau}_{n} \pm z_{\alpha / 2} \widehat{\operatorname{se}}\left(\widehat{\theta}_{n}\right)\left|g^{\prime}\left(\widehat{\theta}_{n}\right)\right| .
$$

Again, this is typically a pointwise asymptotic confidence interval.

### 1.4 Useful Inequalities

At various times in this book we will need to use certain inequalities. For reference purposes, a number of these inequalities are recorded here.

Markov's Inequality. Let $X$ be a non-negative random variable and suppose that $\mathbb{E}(X)$ exists. For any $t>0$,

$$
\begin{equation*}
\mathbb{P}(X>t) \leq \frac{\mathbb{E}(X)}{t} \tag{1.21}
\end{equation*}
$$

Chebyshev's Inequality. Let $\mu=\mathbb{E}(X)$ and $\sigma^{2}=\mathbb{V}(X)$. Then,

$$
\begin{equation*}
\mathbb{P}(|X-\mu| \geq t) \leq \frac{\sigma^{2}}{t^{2}} \tag{1.22}
\end{equation*}
$$

Hoeffding's Inequality. Let $Y_{1}, \ldots, Y_{n}$ be independent observations such that $\mathbb{E}\left(Y_{i}\right)=0$ and $a_{i} \leq Y_{i} \leq b_{i}$. Let $\epsilon>0$. Then, for any $t>0$,

$$
\begin{equation*}
\mathbb{P}\left(\sum_{i=1}^{n} Y_{i} \geq \epsilon\right) \leq e^{-t \epsilon} \prod_{i=1}^{n} e^{t^{2}\left(b_{i}-a_{i}\right)^{2} / 8} \tag{1.23}
\end{equation*}
$$

Hoeffding's Inequality for Bernoulli Random Variables. Let $X_{1}, \ldots, X_{n} \sim \operatorname{Bernoulli}(p)$. Then, for any $\epsilon>0$,

$$
\begin{equation*}
\mathbb{P}\left(\left|\bar{X}_{n}-p\right|>\epsilon\right) \leq 2 e^{-2 n \epsilon^{2}} \tag{1.24}
\end{equation*}
$$

where $\bar{X}_{n}=n^{-1} \sum_{i=1}^{n} X_{i}$.

Mill's Inequality. If $Z \sim N(0,1)$ then, for any $t>0$,

$$
\begin{equation*}
\mathbb{P}(|Z|>t) \leq \frac{2 \phi(t)}{t} \tag{1.25}
\end{equation*}
$$

where $\phi$ is the standard Normal density. In fact, for any $t>0$,

$$
\begin{equation*}
\left(\frac{1}{t}-\frac{1}{t^{3}}\right) \phi(t)<\mathbb{P}(Z>t)<\frac{1}{t} \phi(t) \tag{1.26}
\end{equation*}
$$

and

$$
\begin{equation*}
P(Z>t)<\frac{1}{2} e^{-t^{2} / 2} \tag{1.27}
\end{equation*}
$$

Berry-Esséen Bound. Let $X_{1}, \ldots, X_{n}$ be IID with finite mean $\mu=\mathbb{E}\left(X_{1}\right)$, variance $\sigma^{2}=\mathbb{V}\left(X_{1}\right)$ and third moment, $\mathbb{E}\left|X_{1}\right|^{3}<\infty$. Let $Z_{n}=\sqrt{n}\left(\bar{X}_{n}-\right.$ $\mu) / \sigma$. Then

$$
\begin{equation*}
\sup _{z}\left|\mathbb{P}\left(Z_{n} \leq z\right)-\Phi(z)\right| \leq \frac{33}{4} \frac{\mathbb{E}\left|X_{1}-\mu\right|^{3}}{\sqrt{n} \sigma^{3}} \tag{1.28}
\end{equation*}
$$

Bernstein's Inequality. Let $X_{1}, \ldots, X_{n}$ be independent, zero mean random variables such that $-M \leq X_{i} \leq M$. Then

$$
\begin{equation*}
\mathbb{P}\left(\left|\sum_{i=1}^{n} X_{i}\right|>t\right) \leq 2 \exp \left\{-\frac{1}{2}\left(\frac{t^{2}}{v+M t / 3}\right)\right\} \tag{1.29}
\end{equation*}
$$

where $v \geq \sum_{i=1}^{n} \mathbb{V}\left(X_{i}\right)$.

Bernstein's Inequality (Moment version). Let $X_{1}, \ldots, X_{n}$ be independent, zero mean random variables such that

$$
\mathbb{E}\left|X_{i}\right|^{m} \leq \frac{m!M^{m-2} v_{i}}{2}
$$

for all $m \geq 2$ and some constants $M$ and $v_{i}$. Then,

$$
\begin{equation*}
\mathbb{P}\left(\left|\sum_{i=1}^{n} X_{i}\right|>t\right) \leq 2 \exp \left\{-\frac{1}{2}\left(\frac{t^{2}}{v+M t}\right)\right\} \tag{1.30}
\end{equation*}
$$

where $v=\sum_{i=1}^{n} v_{i}$.

Cauchy-Schwartz Inequality. If $X$ and $Y$ have finite variances then

$$
\begin{equation*}
\mathbb{E}|X Y| \leq \sqrt{\mathbb{E}\left(X^{2}\right) \mathbb{E}\left(Y^{2}\right)} \tag{1.31}
\end{equation*}
$$

Recall that a function $g$ is convex if for each $x, y$ and each $\alpha \in[0,1]$,

$$
g(\alpha x+(1-\alpha) y) \leq \alpha g(x)+(1-\alpha) g(y)
$$

If $g$ is twice differentiable, then convexity reduces to checking that $g^{\prime \prime}(x) \geq 0$ for all $x$. It can be shown that if $g$ is convex then it lies above any line that touches $g$ at some point, called a tangent line. A function $g$ is concave if $-g$ is convex. Examples of convex functions are $g(x)=x^{2}$ and $g(x)=e^{x}$. Examples of concave functions are $g(x)=-x^{2}$ and $g(x)=\log x$.

Jensen's inequality. If $g$ is convex then

$$
\begin{equation*}
\mathbb{E} g(X) \geq g(\mathbb{E} X) \tag{1.32}
\end{equation*}
$$

If $g$ is concave then

$$
\begin{equation*}
\mathbb{E} g(X) \leq g(\mathbb{E} X) \tag{1.33}
\end{equation*}
$$

### 1.5 Bibliographic Remarks

References on probability inequalities and their use in statistics and pattern recognition include Devroye et al. (1996) and van der Vaart and Wellner (1996). To review basic probability and mathematical statistics, I recommend Casella and Berger (2002), van der Vaart (1998) and Wasserman (2004).

### 1.6 Exercises

1. Consider Example 1.17. Prove that (1.18) is a pointwise asymptotic confidence interval. Prove that (1.19) is a uniform confidence interval.
2. (Computer experiment). Compare the coverage and length of (1.18) and (1.19) by simulation. Take $p=0.2$ and use $\alpha=.05$. Try various sample sizes $n$. How large must $n$ be before the pointwise interval has accurate coverage? How do the lengths of the two intervals compare when this sample size is reached?
3. Let $X_{1}, \ldots, X_{n} \sim N(\mu, 1)$. Let $C_{n}=\bar{X}_{n} \pm z_{\alpha / 2} / \sqrt{n}$. Is $C_{n}$ a finite sample, pointwise asymptotic, or uniform asymptotic confidence set for $\mu$ ?
4. Let $X_{1}, \ldots, X_{n} \sim N\left(\mu, \sigma^{2}\right)$. Let $C_{n}=\bar{X}_{n} \pm z_{\alpha / 2} S_{n} / \sqrt{n}$ where $S_{n}^{2}=$ $\sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2} /(n-1)$. Is $C_{n}$ a finite sample, pointwise asymptotic, or uniform asymptotic confidence set for $\mu$ ?
5. Let $X_{1}, \ldots, X_{n} \sim F$ and let $\mu=\int x d F(x)$ be the mean. Let

$$
C_{n}=\left(\bar{X}_{n}-z_{\alpha / 2} \widehat{\operatorname{se}}, \bar{X}_{n}+z_{\alpha / 2} \widehat{\operatorname{se}}\right)
$$

where $\widehat{\operatorname{se}}^{2}=S_{n}^{2} / n$ and

$$
S_{n}^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}
$$

(a) Assuming that the mean exists, show that $C_{n}$ is a $1-\alpha$ pointwise asymptotic confidence interval.
(b) Show that $C_{n}$ is not a uniform asymptotic confidence interval. Hint: Let $a_{n} \rightarrow \infty$ and $\epsilon_{n} \rightarrow 0$ and let $G_{n}=\left(1-\epsilon_{n}\right) F+\epsilon_{n} \delta_{n}$ where $\delta_{n}$ is a pointmass at $a_{n}$. Argue that, with very high probability, for $a_{n}$ large and $\epsilon_{n}$ small, $\int x d G_{n}(x)$ is large but $\bar{X}_{n}+z_{\alpha / 2} \widehat{\text { se }}$ is not large.
(c) Suppose that $\mathbb{P}\left(\left|X_{i}\right| \leq B\right)=1$ where $B$ is a known constant. Use Bernstein's inequality (1.29) to construct a finite sample confidence interval for $\mu$.

## 2

## Estimating the CDF and Statistical Functionals

The first problem we consider is estimating the CDF. By itself, this is not a very interesting problem. However, it is the first step towards solving more important problems such as estimating statistical functionals.

### 2.1 The CDF

We begin with the problem of estimating a CDF (cumulative distribution function). Let $X_{1}, \ldots, X_{n} \sim F$ where $F(x)=\mathbb{P}(X \leq x)$ is a distribution function on the real line. We estimate $F$ with the empirical distribution function.
2.1 Definition. The empirical distribution function $\widehat{F}_{n}$ is the CDF that puts mass $1 / n$ at each data point $X_{i}$. Formally,

$$
\begin{equation*}
\widehat{F}_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} I\left(X_{i} \leq x\right) \tag{2.2}
\end{equation*}
$$

where

$$
I\left(X_{i} \leq x\right)= \begin{cases}1 & \text { if } X_{i} \leq x \\ 0 & \text { if } X_{i}>x\end{cases}
$$



FIGURE 2.1. Nerve data. Each vertical line represents one data point. The solid line is the empirical distribution function. The lines above and below the middle line are a 95 percent confidence band.
2.3 Example (Nerve data). Cox and Lewis (1966) reported 799 waiting times between successive pulses along a nerve fiber. Figure 2.1 shows the data and the empirical CDF $\widehat{F}_{n}$.

The following theorem gives some properties of $\widehat{F}_{n}(x)$.
2.4 Theorem. Let $X_{1}, \ldots, X_{n} \sim F$ and let $\widehat{F}_{n}$ be the empirical CDF. Then:

1. At any fixed value of $x$,

$$
\mathbb{E}\left(\widehat{F}_{n}(x)\right)=F(x) \quad \text { and } \quad \mathbb{V}\left(\widehat{F}_{n}(x)\right)=\frac{F(x)(1-F(x))}{n}
$$

Thus, MSE $=\frac{F(x)(1-F(x))}{n} \rightarrow 0$ and hence $\widehat{F}_{n}(x) \xrightarrow{\mathrm{P}} F(x)$.
2. (Glivenko-Cantelli Theorem).

$$
\sup _{x}\left|\widehat{F}_{n}(x)-F(x)\right| \xrightarrow{\text { a.s. }} 0 .
$$

3. (Dvoretzky-Kiefer-Wolfowitz (DKW) inequality). For any $\epsilon>0$,

$$
\begin{equation*}
\mathbb{P}\left(\sup _{x}\left|F(x)-\widehat{F}_{n}(x)\right|>\epsilon\right) \leq 2 e^{-2 n \epsilon^{2}} . \tag{2.5}
\end{equation*}
$$

From the DKW inequality, we can construct a confidence set. Let $\epsilon_{n}^{2}=$ $\log (2 / \alpha) /(2 n), L(x)=\max \left\{\widehat{F}_{n}(x)-\epsilon_{n}, 0\right\}$ and $U(x)=\min \left\{\widehat{F}_{n}(x)+\epsilon_{n}, 1\right\}$. It follows from (2.5) that for any $F$,

$$
\mathbb{P}(L(x) \leq F(x) \leq U(x) \text { for all } x) \geq 1-\alpha
$$

Thus, $(L(x), U(x))$ is a nonparametric $1-\alpha$ confidence band. ${ }^{1}$ To summarize:
2.6 Theorem. Let

$$
\begin{aligned}
L(x) & =\max \left\{\widehat{F}_{n}(x)-\epsilon_{n}, 0\right\} \\
U(x) & =\min \left\{\widehat{F}_{n}(x)+\epsilon_{n}, 1\right\}
\end{aligned}
$$

where

$$
\epsilon_{n}=\sqrt{\frac{1}{2 n} \log \left(\frac{2}{\alpha}\right)}
$$

Then, for all $F$ and all $n$,

$$
\mathbb{P}(L(x) \leq F(x) \leq U(x) \text { for all } x) \geq 1-\alpha
$$

2.7 Example. The dashed lines in Figure 2.1 give a 95 percent confidence band using $\epsilon_{n}=\sqrt{\frac{1}{2 n} \log \left(\frac{2}{.05}\right)}=.048$.

### 2.2 Estimating Statistical Functionals

A statistical functional $T(F)$ is any function of $F$. Examples are the mean $\mu=\int x d F(x)$, the variance $\sigma^{2}=\int(x-\mu)^{2} d F(x)$ and the median $m=$ $F^{-1}(1 / 2)$.
2.8 Definition. The plug-in estimator of $\theta=T(F)$ is defined by

$$
\begin{equation*}
\widehat{\theta}_{n}=T\left(\widehat{F}_{n}\right) . \tag{2.9}
\end{equation*}
$$

A functional of the form $\int a(x) d F(x)$ is called a linear functional. Recall that $\int a(x) d F(x)$ is defined to be $\int a(x) f(x) d x$ in the continuous case

[^0]and $\sum_{j} a\left(x_{j}\right) f\left(x_{j}\right)$ in the discrete case. The empirical CDF $\widehat{F}_{n}(x)$ is discrete, putting mass $1 / n$ at each $X_{i}$. Hence, if $T(F)=\int a(x) d F(x)$ is a linear functional then we have:

The plug-in estimator for linear functional $T(F)=\int a(x) d F(x)$ is:

$$
\begin{equation*}
T\left(\widehat{F}_{n}\right)=\int a(x) d \widehat{F}_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} a\left(X_{i}\right) . \tag{2.10}
\end{equation*}
$$

Sometimes we can find the estimated standard error se of $T\left(\widehat{F}_{n}\right)$ by doing some direct calculations. However, in other cases it is not obvious how to estimate the standard error. Later we will discuss methods for finding se. For now, let us just assume that somehow we can find se. In many cases, it turns out that

$$
\begin{equation*}
T\left(\widehat{F}_{n}\right) \approx N\left(T(F),{\widehat{\mathrm{se}^{2}}}^{2}\right) \tag{2.11}
\end{equation*}
$$

In that case, an approximate $1-\alpha$ confidence interval for $T(F)$ is then

$$
\begin{equation*}
T\left(\widehat{F}_{n}\right) \pm z_{\alpha / 2} \widehat{\mathrm{se}} \tag{2.12}
\end{equation*}
$$

where $z_{\alpha}$ is defined by: $\mathbb{P}\left(Z>z_{\alpha}\right)=\alpha$ with $Z \sim N(0,1)$. We will call (2.12) the Normal-based interval.
2.13 Example (The mean). Let $\mu=T(F)=\int x d F(x)$. The plug-in estimator is $\widehat{\mu}=\int x d \widehat{F}_{n}(x)=\bar{X}_{n}$. The standard error is se $=\sqrt{\mathbb{V}\left(\bar{X}_{n}\right)}=\sigma / \sqrt{n}$. If $\widehat{\sigma}$ denotes an estimate of $\sigma$, then the estimated standard error is $\widehat{\operatorname{se}}=\widehat{\sigma} / \sqrt{n}$. A Normal-based confidence interval for $\mu$ is $\bar{X}_{n} \pm z_{\alpha / 2} \widehat{\sigma} / \sqrt{n}$.
2.14 Example (The variance). Let $\sigma^{2}=\mathbb{V}(X)=\int x^{2} d F(x)-\left(\int x d F(x)\right)^{2}$. The plug-in estimator is

$$
\begin{aligned}
\widehat{\sigma}^{2} & =\int x^{2} d \widehat{F}_{n}(x)-\left(\int x d \widehat{F}_{n}(x)\right)^{2} \\
& =\frac{1}{n} \sum_{i=1}^{n} X_{i}^{2}-\left(\frac{1}{n} \sum_{i=1}^{n} X_{i}\right)^{2} \\
& =\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}
\end{aligned}
$$

This is different than the usual unbiased sample variance

$$
S_{n}^{2}=\frac{1}{n-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2} .
$$

In practice, there is little difference between $\widehat{\sigma}^{2}$ and $S_{n}^{2}$.
2.15 Example (The skewness). Let $\mu$ and $\sigma^{2}$ denote the mean and variance of a random variable $X$. The skewness - which measures the lack of symmetry of a distribution - is defined to be

$$
\kappa=\frac{\mathbb{E}(X-\mu)^{3}}{\sigma^{3}}=\frac{\int(x-\mu)^{3} d F(x)}{\left\{\int(x-\mu)^{2} d F(x)\right\}^{3 / 2}}
$$

To find the plug-in estimate, first recall that $\widehat{\mu}=n^{-1} \sum_{i=1}^{n} X_{i}$ and $\widehat{\sigma}^{2}=$ $n^{-1} \sum_{i=1}^{n}\left(X_{i}-\widehat{\mu}\right)^{2}$. The plug-in estimate of $\kappa$ is

$$
\widehat{\kappa}=\frac{\int(x-\mu)^{3} d \widehat{F}_{n}(x)}{\left\{\int(x-\mu)^{2} d \widehat{F}_{n}(x)\right\}^{3 / 2}}=\frac{\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\widehat{\mu}\right)^{3}}{\widehat{\sigma}^{3}}
$$

2.16 Example (Correlation). Let $Z=(X, Y)$ and let $\rho=T(F)=\mathbb{E}(X-$ $\left.\mu_{X}\right)\left(Y-\mu_{Y}\right) /\left(\sigma_{x} \sigma_{y}\right)$ denote the correlation between $X$ and $Y$, where $F(x, y)$ is bivariate. We can write $T(F)=a\left(T_{1}(F), T_{2}(F), T_{3}(F), T_{4}(F), T_{5}(F)\right)$ where

$$
\begin{array}{lll}
T_{1}(F)=\int x d F(z) & T_{2}(F)=\int y d F(z) & T_{3}(F)=\int x y d F(z) \\
T_{4}(F)=\int x^{2} d F(z) & T_{5}(F)=\int y^{2} d F(z)
\end{array}
$$

and

$$
a\left(t_{1}, \ldots, t_{5}\right)=\frac{t_{3}-t_{1} t_{2}}{\sqrt{\left(t_{4}-t_{1}^{2}\right)\left(t_{5}-t_{2}^{2}\right)}}
$$

Replace $F$ with $\widehat{F}_{n}$ in $T_{1}(F), \ldots, T_{5}(F)$, and take

$$
\widehat{\rho}=a\left(T_{1}\left(\widehat{F}_{n}\right), T_{2}\left(\widehat{F}_{n}\right), T_{3}\left(\widehat{F}_{n}\right), T_{4}\left(\widehat{F}_{n}\right), T_{5}\left(\widehat{F}_{n}\right)\right)
$$

We get

$$
\widehat{\rho}=\frac{\sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)\left(Y_{i}-\bar{Y}_{n}\right)}{\sqrt{\sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}} \sqrt{\sum_{i=1}^{n}\left(Y_{i}-\bar{Y}_{n}\right)^{2}}}
$$

which is called the sample correlation.
2.17 Example (Quantiles). Let $F$ be strictly increasing with density $f$. Let $T(F)=F^{-1}(p)$ be the $p^{\text {th }}$ quantile. The estimate of $T(F)$ is $\widehat{F}_{n}^{-1}(p)$. We have to be a bit careful since $\widehat{F}_{n}$ is not invertible. To avoid ambiguity we define $\widehat{F}_{n}^{-1}(p)=\inf \left\{x: \widehat{F}_{n}(x) \geq p\right\}$. We call $\widehat{F}_{n}^{-1}(p)$ the $p^{t h}$ sample quantile.

The Glivenko-Cantelli theorem ensures that $\widehat{F}_{n}$ converges to $F$. This suggests that $\widehat{\theta}_{n}=T\left(\widehat{F}_{n}\right)$ will converge to $\theta=T(F)$. Furthermore, we would hope that under reasonable conditions, $\widehat{\theta}_{n}$ will be asymptotically normal. This leads us to the next topic.

### 2.3 Influence Functions

The influence function is used to approximate the standard error of a plug-in estimator. The formal definition is as follows.
2.18 Definition. The Gâteaux derivative of $T$ at $F$ in the direction $G$ is defined by

$$
\begin{equation*}
L_{F}(G)=\lim _{\epsilon \rightarrow 0} \frac{T((1-\epsilon) F+\epsilon G)-T(F)}{\epsilon} . \tag{2.19}
\end{equation*}
$$

If $G=\delta_{x}$ is a point mass at $x$ then we write $L_{F}(x) \equiv L_{F}\left(\delta_{x}\right)$ and we call $L_{F}(x)$ the influence function. Thus,

$$
\begin{equation*}
L_{F}(x)=\lim _{\epsilon \rightarrow 0} \frac{T\left((1-\epsilon) F+\epsilon \delta_{x}\right)-T(F)}{\epsilon} . \tag{2.20}
\end{equation*}
$$

The empirical influence function is defined by $\widehat{L}(x)=L_{\widehat{F}_{n}}(x)$. Thus,

$$
\begin{equation*}
\widehat{L}(x)=\lim _{\epsilon \rightarrow 0} \frac{T\left((1-\epsilon) \widehat{F}_{n}+\epsilon \delta_{x}\right)-T\left(\widehat{F}_{n}\right)}{\epsilon} . \tag{2.21}
\end{equation*}
$$

Often we drop the subscript $F$ and write $L(x)$ instead of $L_{F}(x)$.
2.22 Theorem. Let $T(F)=\int a(x) d F(x)$ be a linear functional. Then:

1. $L_{F}(x)=a(x)-T(F)$ and $\widehat{L}(x)=a(x)-T\left(\widehat{F}_{n}\right)$.
2. For any $G$,

$$
\begin{equation*}
T(G)=T(F)+\int L_{F}(x) d G(x) \tag{2.23}
\end{equation*}
$$

3. $\int L_{F}(x) d F(x)=0$.
4. Let $\tau^{2}=\int L_{F}^{2}(x) d F(x)$. Then, $\tau^{2}=\int(a(x)-T(F))^{2} d F(x)$ and if $\tau^{2}<\infty$,

$$
\begin{equation*}
\sqrt{n}\left(T(F)-T\left(\widehat{F}_{n}\right)\right) \rightsquigarrow N\left(0, \tau^{2}\right) . \tag{2.24}
\end{equation*}
$$

5. Let

$$
\begin{equation*}
\widehat{\tau}^{2}=\frac{1}{n} \sum_{i=1}^{n} \widehat{L}^{2}\left(X_{i}\right)=\frac{1}{n} \sum_{i=1}^{n}\left(a\left(X_{i}\right)-T\left(\widehat{F}_{n}\right)\right)^{2} . \tag{2.25}
\end{equation*}
$$

Then, $\widehat{\tau}^{2} \xrightarrow{\mathrm{P}} \tau^{2}$ and $\widehat{\mathrm{se}} / \mathrm{se} \xrightarrow{\mathrm{P}} 1$ where $\widehat{\mathrm{se}}=\widehat{\tau} / \sqrt{n}$ and se $=\sqrt{\mathbb{V}\left(T\left(\widehat{F}_{n}\right)\right)}$.
6. We have that

$$
\begin{equation*}
\frac{\sqrt{n}\left(T(F)-T\left(\widehat{F}_{n}\right)\right)}{\widehat{\tau}} \rightsquigarrow N(0,1) . \tag{2.26}
\end{equation*}
$$

Proof. The first three claims follow easily from the definition of the influence function. To prove the fourth claim, write

$$
\begin{aligned}
T\left(\widehat{F}_{n}\right) & =T(F)+\int L_{F}(x) d \widehat{F}_{n}(x) \\
& =T(F)+\frac{1}{n} \sum_{i=1}^{n} L_{F}\left(X_{i}\right)
\end{aligned}
$$

From the central limit theorem and the fact that $\int L_{F}(x) d F(x)=0$, it follows that

$$
\sqrt{n}\left(T(F)-T\left(\widehat{F}_{n}\right)\right) \rightsquigarrow N\left(0, \tau^{2}\right)
$$

where $\tau^{2}=\int L_{F}^{2}(x) d F(x)$. The fifth claim follows from the law of large numbers. The final statement follows from the fourth and fifth claims and Slutsky's theorem.

The theorem above tells us that the influence function $L_{F}(x)$ behaves like the score function in parametric estimation. To see this, recall that if $f(x ; \theta)$ is a parametric model, $\mathcal{L}_{n}(\theta)=\prod_{\hat{\theta}=1}^{n} f\left(X_{i} ; \theta\right)$ is the likelihood function and the maximum likelihood estimator $\widehat{\theta}_{n}$ is the value of $\theta$ that maximizes $\mathcal{L}_{n}(\theta)$. The score function is $s_{\theta}(x)=\partial \log f(x ; \theta) / \partial \theta$ which, under appropriate regularity conditions, satisfies $\int s_{\theta}(x) f(x ; \theta) d x=0$ and $\mathbb{V}\left(\widehat{\theta}_{n}\right) \approx \int\left(s_{\theta}(x)\right)^{2} f(x ; \theta) d x / n$. Similarly, for the influence function we have that $\int L_{F}(x) d F(x)=0$ and and $\mathbb{V}\left(T\left(\widehat{F}_{n}\right)\right) \approx \int L_{F}^{2}(x) d F(x) / n$.

If the functional $T(F)$ is not linear, then (2.23) will not hold exactly, but it may hold approximately.
2.27 Theorem. If $T$ is Hadamard differentiable ${ }^{2}$ with respect to $d(F, G)=$ $\sup _{x}|F(x)-G(x)|$ then

$$
\begin{equation*}
\sqrt{n}\left(T\left(\widehat{F}_{n}\right)-T(F)\right) \rightsquigarrow N\left(0, \tau^{2}\right) \tag{2.28}
\end{equation*}
$$

where $\tau^{2}=\int L_{F}(x)^{2} d F(x)$. Also,

$$
\begin{equation*}
\frac{\left(T\left(\widehat{F}_{n}\right)-T(F)\right)}{\widehat{\mathrm{se}}} \rightsquigarrow N(0,1) \tag{2.29}
\end{equation*}
$$

[^1]where $\widehat{\text { se }}=\widehat{\tau} / \sqrt{n}$ and
\[

$$
\begin{equation*}
\widehat{\tau}=\frac{1}{n} \sum_{i=1}^{n} L^{2}\left(X_{i}\right) \tag{2.30}
\end{equation*}
$$

\]

We call the approximation $\left(T\left(\widehat{F}_{n}\right)-T(F)\right) / \widehat{\text { se }} \approx N(0,1)$ the nonparametric delta method. From the normal approximation, a large sample confidence interval is $T\left(\widehat{F}_{n}\right) \pm z_{\alpha / 2} \widehat{\text { se }}$. This is only a pointwise asymptotic confidence interval. In summary:

## The Nonparametric Delta Method

A $1-\alpha$, pointwise asymptotic confidence interval for $T(F)$ is

$$
\begin{equation*}
T\left(\widehat{F}_{n}\right) \pm z_{\alpha / 2} \widehat{\text { se }} \tag{2.31}
\end{equation*}
$$

where

$$
\widehat{\mathrm{se}}=\frac{\widehat{\tau}}{\sqrt{n}} \quad \text { and } \quad \widehat{\tau}^{2}=\frac{1}{n} \sum_{i=1}^{n} \widehat{L}^{2}\left(X_{i}\right)
$$

2.32 Example (The mean). Let $\theta=T(F)=\int x d F(x)$. The plug-in estimator is $\widehat{\theta}=\int x d \widehat{F}_{n}(x)=\bar{X}_{n}$. Also, $T\left((1-\epsilon) F+\epsilon \delta_{x}\right)=(1-\epsilon) \theta+\epsilon x$. Thus, $L(x)=x-\theta, \widehat{L}(x)=x-\bar{X}_{n}$ and $\widehat{\mathrm{se}}^{2}=\widehat{\sigma}^{2} / n$ where $\widehat{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}$. A pointwise asymptotic nonparametric 95 percent confidence interval for $\theta$ is $\bar{X}_{n} \pm 2 \widehat{\text { se. }}$

Sometimes statistical functionals take the form $T(F)=a\left(T_{1}(F), \ldots, T_{m}(F)\right)$ for some function $a\left(t_{1}, \ldots, t_{m}\right)$. By the chain rule, the influence function is

$$
L(x)=\sum_{i=1}^{m} \frac{\partial a}{\partial t_{i}} L_{i}(x)
$$

where

$$
\begin{equation*}
L_{i}(x)=\lim _{\epsilon \rightarrow 0} \frac{T_{i}\left((1-\epsilon) F+\epsilon \delta_{x}\right)-T_{i}(F)}{\epsilon} . \tag{2.33}
\end{equation*}
$$

2.34 Example (Correlation). Let $Z=(X, Y)$ and let $T(F)=E\left(X-\mu_{X}\right)(Y-$ $\left.\mu_{Y}\right) /\left(\sigma_{x} \sigma_{y}\right)$ denote the correlation where $F(x, y)$ is bivariate. Recall that $T(F)=a\left(T_{1}(F), T_{2}(F), T_{3}(F), T_{4}(F), T_{5}(F)\right)$ where

$$
\begin{array}{lll}
T_{1}(F)=\int x d F(z) & T_{2}(F)=\int y d F(z) & T_{3}(F)=\int x y d F(z) \\
T_{4}(F)=\int x^{2} d F(z) & T_{5}(F)=\int y^{2} d F(z)
\end{array}
$$

and

$$
a\left(t_{1}, \ldots, t_{5}\right)=\frac{t_{3}-t_{1} t_{2}}{\sqrt{\left(t_{4}-t_{1}^{2}\right)\left(t_{5}-t_{2}^{2}\right)}}
$$

It follows from (2.33) that

$$
L(x, y)=\widetilde{x} \widetilde{y}-\frac{1}{2} T(F)\left(\widetilde{x}^{2}+\widetilde{y}^{2}\right)
$$

where

$$
\widetilde{x}=\frac{x-\int x d F}{\sqrt{\int x^{2} d F-\left(\int x d F\right)^{2}}}, \quad \widetilde{y}=\frac{y-\int y d F}{\sqrt{\int y^{2} d F-\left(\int y d F\right)^{2}}} .
$$

2.35 Example (Quantiles). Let $F$ be strictly increasing with positive density $f$. The $T(F)=F^{-1}(p)$ be the $p^{\text {th }}$ quantile. The influence function is (see Exercise 10)

$$
L(x)= \begin{cases}\frac{p-1}{f(\theta)}, & x \leq \theta \\ \frac{p}{f(\theta)}, & x>\theta\end{cases}
$$

The asymptotic variance of $T\left(\widehat{F}_{n}\right)$ is

$$
\begin{equation*}
\frac{\tau^{2}}{n}=\frac{1}{n} \int L^{2}(x) d F(x)=\frac{p(1-p)}{n f^{2}(\theta)} \tag{2.36}
\end{equation*}
$$

To estimate this variance we need to estimate the density $f$. Later we shall see that the bootstrap provides a simpler estimate of the variance.

### 2.4 Empirical Probability Distributions

This section discusses a generalization of the DKW inequality. The reader may skip this section if desired. Using the empirical CDF to estimate the true CDF is a special case of a more general idea. Let $X_{1}, \ldots, X_{n} \sim P$ be an IID sample from a probability measure $P$. Define the empirical probability distribution $\mathbb{P}_{n}$ by

$$
\begin{equation*}
\widehat{\mathbb{P}}_{n}(A)=\frac{\text { number of } X_{i} \in A}{n} . \tag{2.37}
\end{equation*}
$$

We would like to be able to say that $\mathbb{P}_{n}$ is close to $P$ in some sense. For a fixed $A$ we know that $n \widehat{P}_{n}(A) \sim \operatorname{Binomial}(n, p)$ where $p=P(A)$. By Hoeffding's inequality, it follows that

$$
\begin{equation*}
\mathbb{P}\left(\left|\widehat{P}_{n}(A)-P(A)\right|>\epsilon\right) \leq 2 e^{-2 n \epsilon^{2}} \tag{2.38}
\end{equation*}
$$

We would like to extend this to be a statement of the form

$$
\mathbb{P}\left(\sup _{A \in \mathcal{A}}\left|\widehat{P}_{n}(A)-P(A)\right|>\epsilon\right) \leq \text { something small }
$$

for some class of sets $\mathcal{A}$. This is exactly what the DKW inequality does by taking $\mathcal{A}=\{A=(-\infty, t]: t \in \mathbb{R}\}$. But DKW is only useful for onedimensional random variables. We can get a more general inequality by using Vapnik-Chervonenkis (VC) theory.

Let $\mathcal{A}$ be a class of sets. Given a finite set $R=\left\{x_{1}, \ldots, x_{n}\right\}$ let

$$
\begin{equation*}
N_{\mathcal{A}}(R)=\#\{R \bigcap A: A \in \mathcal{A}\} \tag{2.39}
\end{equation*}
$$

be the number of subsets of $R$ "picked out" as $A$ varies over $\mathcal{A}$. We say that $R$ is shattered by $\mathcal{A}$ if $N_{\mathcal{A}}(R)=2^{n}$. The shatter coefficient is defined by

$$
\begin{equation*}
s(\mathcal{A}, n)=\max _{R \in \mathcal{F}_{n}} N_{\mathcal{A}}(R) \tag{2.40}
\end{equation*}
$$

where $\mathcal{F}_{n}$ consists of all finite sets of size $n$.
2.41 Theorem (Vapnik and Chervonenkis, 1971). For any $P$, $n$ and $\epsilon>0$,

$$
\begin{equation*}
\mathbb{P}\left(\sup _{A \in \mathcal{A}}\left|\widehat{P}_{n}(A)-P(A)\right|>\epsilon\right) \leq 8 s(\mathcal{A}, n) e^{-n \epsilon^{2} / 32} \tag{2.42}
\end{equation*}
$$

Theorem 2.41 is only useful if the shatter coefficients do not grow too quickly with $n$. This is where VC dimension enters. If $s(\mathcal{A}, n)=2^{n}$ for all $n$ set $\mathrm{VC}(\mathcal{A})=\infty$. Otherwise, define $\operatorname{VC}(\mathcal{A})$ to be the largest $k$ for which $s(\mathcal{A}, k)=$ $2^{k}$. We call $\operatorname{VC}(\mathcal{A})$ the Vapnik-Chervonenkis dimension of $\mathcal{A}$. Thus, the VC-dimension is the size of the largest finite set $F$ that is shattered by $\mathcal{A}$. The following theorem shows that if $\mathcal{A}$ has finite VC-dimension then the shatter coefficients grow as a polynomial in $n$.
2.43 Theorem. If $\mathcal{A}$ has finite $V C$-dimension $v$, then

$$
s(\mathcal{A}, n) \leq n^{v}+1
$$

In this case,

$$
\begin{equation*}
\mathbb{P}\left(\sup _{A \in \mathcal{A}}\left|\widehat{P}_{n}(A)-P(A)\right|>\epsilon\right) \leq 8\left(n^{v}+1\right) e^{-n \epsilon^{2} / 32} \tag{2.44}
\end{equation*}
$$

2.45 Example. Let $\mathcal{A}=\{(-\infty, x] ; x \in \mathbb{R}\}$. Then $\mathcal{A}$ shatters every one point set $\{x\}$ but it shatters no set of the form $\{x, y\}$. Therefore, $\operatorname{VC}(\mathcal{A})=1$. Since, $\mathbb{P}((-\infty, x])=F(x)$ is the CDF and $\widehat{\mathbb{P}}_{n}((-\infty, x])=\widehat{F}_{n}(x)$ is the empirical CDF, we conclude that

$$
\mathbb{P}\left(\sup _{x}\left|\widehat{F}_{n}(x)-F(x)\right|>\epsilon\right) \leq 8(n+1) e^{-n \epsilon^{2} / 32}
$$

which is looser than the DKW bound. This shows that the bound (2.42) is not the tightest possible.
2.46 Example. Let $\mathcal{A}$ be the set of closed intervals on the real line. Then $\mathcal{A}$ shatters $S=\{x, y\}$ but it cannot shatter sets with three points. Consider $S=\{x, y, z\}$ where $x<y<z$. One cannot find an interval $A$ such that $A \bigcap S=\{x, z\}$. So, $\operatorname{VC}(\mathcal{A})=2$.
2.47 Example. Let $\mathcal{A}$ be all linear half-spaces on the plane. Any three-point set (not all on a line) can be shattered. No four-point set can be shattered. Consider, for example, four points forming a diamond. Let $T$ be the leftmost and rightmost points. This set cannot be picked out. Other configurations can also be seen to be unshatterable. So $\operatorname{VC}(\mathcal{A})=3$. In general, halfspaces in $\mathcal{R}^{d}$ have VC dimension $d+1$.
2.48 Example. Let $\mathcal{A}$ be all rectangles on the plane with sides parallel to the axes. Any four-point set can be shattered. Let $S$ be a five-point set. There is one point that is not leftmost, rightmost, uppermost or lowermost. Let $T$ be all points in $S$ except this point. Then $T$ can't be picked out. So, we have that $\operatorname{VC}(\mathcal{A})=4$.

### 2.5 Bibliographic Remarks

Further details on statistical functionals can be found in Serfling (1980), Davison and Hinkley (1997), Shao and Tu (1995), Fernholz (1983) and van der Vaart (1998). Vapnik-Chervonenkis theory is discussed in Devroye et al. (1996), van der Vaart (1998) and van der Vaart and Wellner (1996).

### 2.6 Appendix

Here are some details about Theorem 2.27. Let $\mathfrak{F}$ denote all distribution functions and let $\mathcal{D}$ denote the linear space generated by $\mathfrak{F}$. Write $T((1-\epsilon) F+$ $\epsilon G)=T(F+\epsilon D)$ where $D=G-F \in \mathcal{D}$. The Gateâux derivative, which we now write as $L_{F}(D)$, is defined by

$$
\lim _{\epsilon \downarrow 0}\left|\frac{T(F+\epsilon D)-T(F)}{\epsilon}-L_{F}(D)\right| \rightarrow 0 .
$$

Thus $T(F+\epsilon D) \approx \epsilon L_{F}(D)+o(\epsilon)$ and the error term $o(\epsilon)$ goes to 0 as $\epsilon \rightarrow 0$. Hadamard differentiability requires that this error term be small uniformly over compact sets. Equip $\mathcal{D}$ with a metric $d . T$ is Hadamard differentiable
at $F$ if there exists a linear functional $L_{F}$ on $\mathcal{D}$ such that for any $\epsilon_{n} \rightarrow 0$ and $\left\{D, D_{1}, D_{2}, \ldots\right\} \subset \mathcal{D}$ such that $d\left(D_{n}, D\right) \rightarrow 0$ and $F+\epsilon_{n} D_{n} \in \mathcal{F}$,

$$
\lim _{n \rightarrow \infty}\left(\frac{T\left(F+\epsilon_{n} D_{n}\right)-T(F)}{\epsilon_{n}}-L_{F}\left(D_{n}\right)\right)=0
$$

### 2.7 Exercises

1. Fill in the details of the proof of Theorem 2.22.
2. Prove Theorem 2.4.
3. (Computer experiment.) Generate 100 observations from a $N(0,1)$ distribution. Compute a 95 percent confidence band for the CDF $F$. Repeat this 1000 times and see how often the confidence band contains the true distribution function. Repeat using data from a Cauchy distribution.
4. Let $X_{1}, \ldots, X_{n} \sim F$ and let $\widehat{F}_{n}(x)$ be the empirical distribution function. For a fixed $x$, find the limiting distribution of $\sqrt{\widehat{F}_{n}(x)}$.
5. Suppose that

$$
\begin{equation*}
|T(F)-T(G)| \leq C\|F-G\|_{\infty} \tag{2.49}
\end{equation*}
$$

for some constant $0<C<\infty$ where $\|F-G\|_{\infty}=\sup _{x}|F(x)-G(x)|$. Prove that $T\left(\widehat{F}_{n}\right) \xrightarrow{\text { a.s. }} T(F)$. Suppose that $|X| \leq M<\infty$. Show that $T(F)=\int x d F(x)$ satisfies (2.49).

6 . Let $x$ and $y$ be two distinct points. Find $\operatorname{Cov}\left(\widehat{F}_{n}(x), \widehat{F}_{n}(y)\right)$.
7. Let $X_{1}, \ldots, X_{n} \sim \operatorname{Bernoulli}(p)$ and let $Y_{1}, \ldots, Y_{m} \sim \operatorname{Bernoulli}(q)$. Find the plug-in estimator and estimated standard error for $p$. Find an approximate 90 percent confidence interval for $p$. Find the plug-in estimator and estimated standard error for $p-q$. Find an approximate 90 percent confidence interval for $p-q$.
8. Let $X_{1}, \ldots, X_{n} \sim F$ and let $\widehat{F}$ be the empirical distribution function. Let $a<b$ be fixed numbers and define $\theta=T(F)=F(b)-F(a)$. Let $\widehat{\theta}=T\left(\widehat{F}_{n}\right)=\widehat{F}_{n}(b)-\widehat{F}_{n}(a)$. Find the influence function. Find the estimated standard error of $\widehat{\theta}$. Find an expression for an approximate $1-\alpha$ confidence interval for $\theta$.
9. Verify the formula for the influence function in Example 2.34.
10. Verify the formula for the influence function in Example 2.35. Hint: Let $F_{\epsilon}(y)=(1-\epsilon) F(y)+\epsilon \delta_{x}(y)$ where $\delta_{x}(y)$ is a point mass at $x$, i.e., $\delta_{x}(y)=0$ if $y<x$ and $\delta_{x}(y)=1$ if $y \geq x$. By the definition of $T(F)$, we have that $p=F_{\epsilon}\left(T\left(F_{\epsilon}\right)\right)$. Now differentiate with respect to $\epsilon$ and evaluate the derivative at $\epsilon=0$.
11. Data on the magnitudes of earthquakes near Fiji are available on the book website. Estimate the CDF $F(x)$. Compute and plot a 95 percent confidence envelope for $F$. Find an approximate 95 percent confidence interval for $F(4.9)-F(4.3)$.
12. Get the data on eruption times and waiting times between eruptions of the Old Faithful geyser from the book website. Estimate the mean waiting time and give a standard error for the estimate. Also, give a 90 percent confidence interval for the mean waiting time. Now estimate the median waiting time. In the next chapter we will see how to get the standard error for the median.
13. In 1975 , an experiment was conducted to see if cloud seeding produced rainfall. 26 clouds were seeded with silver nitrate and 26 were not. The decision to seed or not was made at random. Get the data from
http://lib.stat.cmu.edu/DASL/Stories/CloudSeeding.html
Let $\theta=T\left(F_{1}\right)-T\left(F_{2}\right)$ be the difference in the median precipitation from the two groups. Estimate $\theta$. Estimate the standard error of the estimate and produce a 95 percent confidence interval. To estimate the standard error you will need to use formula (2.36). This formula requires the density $f$ so you will need to insert an estimate of $f$. What will you do? Be creative.
14. Let $\mathcal{A}$ be the set of two-dimensional spheres. That is, $A \in \mathcal{A}$ if $A=$ $\left\{(x, y):(x-a)^{2}+(y-b)^{2} \leq c^{2}\right\}$ for some $a, b, c$. Find the VC dimension of $\mathcal{A}$.
15. The empirical CDF can be regarded as a nonparametric maximum likelihood estimator. For example, consider data $X_{1}, \ldots, X_{n}$ on $[0,1]$. Divide the interval into bins of width $\Delta$ and find the MLE over all distributions with constant density over the bins. Show that the resulting CDF converges to the empirical CDF as $\Delta \rightarrow 0$.

## 3

## The Bootstrap and the Jackknife

The bootstrap and the jackknife are nonparametric methods for computing standard errors and confidence intervals. The jackknife is less computationally expensive, but the bootstrap has some statistical advantages.

### 3.1 The Jackknife

The jackknife, due to Quenouille (1949), is a simple method for approximating the bias and variance of an estimator. Let $T_{n}=T\left(X_{1}, \ldots, X_{n}\right)$ be an estimator of some quantity $\theta$ and let $\operatorname{bias}\left(T_{n}\right)=\mathbb{E}\left(T_{n}\right)-\theta$ denote the bias. Let $T_{(-i)}$ denote the statistic with the $i^{\text {th }}$ observation removed. The jackknife bias estimate is defined by

$$
\begin{equation*}
b_{\mathrm{jack}}=(n-1)\left(\bar{T}_{n}-T_{n}\right) \tag{3.1}
\end{equation*}
$$

where $\bar{T}_{n}=n^{-1} \sum_{i} T_{(-i)}$. The bias-corrected estimator is $T_{\mathrm{jack}}=T_{n}-b_{\mathrm{jack}}$.
Why is $b_{\text {jack }}$ defined this way? For many statistics it can be shown that

$$
\begin{equation*}
\operatorname{bias}\left(T_{n}\right)=\frac{a}{n}+\frac{b}{n^{2}}+O\left(\frac{1}{n^{3}}\right) \tag{3.2}
\end{equation*}
$$

for some $a$ and $b$. For example, let $\sigma^{2}=\mathbb{V}\left(X_{i}\right)$ and let $\widehat{\sigma}_{n}^{2}=n^{-1} \sum_{i=1}^{n}\left(X_{i}-\right.$ $\bar{X})^{2}$. Then, $\mathbb{E}\left(\widehat{\sigma}_{n}^{2}\right)=(n-1) \sigma^{2} / n$ so that $\operatorname{bias}\left(\widehat{\sigma}_{n}^{2}\right)=-\sigma^{2} / n$. Thus, (3.2) holds with $a=-\sigma^{2}$ and $b=0$.

When (3.2) holds, we have

$$
\begin{equation*}
\operatorname{bias}\left(T_{(-i)}\right)=\frac{a}{n-1}+\frac{b}{(n-1)^{2}}+O\left(\frac{1}{n^{3}}\right) \tag{3.3}
\end{equation*}
$$

It follows that $\operatorname{bias}\left(\bar{T}_{n}\right)$ also satisfies (3.3). Hence,

$$
\begin{aligned}
\mathbb{E}\left(b_{\mathrm{jack}}\right) & =(n-1)\left(\mathbb{E}\left(\operatorname{bias}\left(\bar{T}_{n}\right)\right)-\mathbb{E}\left(\operatorname{bias}\left(T_{n}\right)\right)\right) \\
& =(n-1)\left[\left(\frac{1}{n-1}-\frac{1}{n}\right) a+\left(\frac{1}{(n-1)^{2}}-\frac{1}{n^{2}}\right) b+O\left(\frac{1}{n^{3}}\right)\right] \\
& =\frac{a}{n}+\frac{(2 n-1) b}{n^{2}(n-1)}+O\left(\frac{1}{n^{2}}\right) \\
& =\operatorname{bias}\left(T_{n}\right)+O\left(\frac{1}{n^{2}}\right)
\end{aligned}
$$

which shows that $b_{\text {jack }}$ estimates the bias up to order $O\left(n^{-2}\right)$. By a similar calculation,

$$
\operatorname{bias}\left(T_{\mathrm{jack}}\right)=-\frac{b}{n(n-1)}+O\left(\frac{1}{n^{2}}\right)=O\left(\frac{1}{n^{2}}\right)
$$

so the bias of $T_{\text {jack }}$ is an order of magnitude smaller than that of $T_{n} . T_{\text {jack }}$ can also be written as

$$
T_{\mathrm{jack}}=\frac{1}{n} \sum_{i=1}^{n} \widetilde{T}_{i}
$$

where

$$
\widetilde{T}_{i}=n T_{n}-(n-1) T_{(-i)}
$$

are called pseudo-values.
The jackknife estimate of $\mathbb{V}\left(T_{n}\right)$ is

$$
\begin{equation*}
v_{\mathrm{jack}}=\frac{\widetilde{s}^{2}}{n} \tag{3.4}
\end{equation*}
$$

where

$$
\widetilde{s}^{2}=\frac{\sum_{i=1}^{n}\left(\widetilde{T}_{i}-\frac{1}{n} \sum_{i=1}^{n} \widetilde{T}_{i}\right)^{2}}{n-1}
$$

is the sample variance of the pseudo-values. Under suitable conditions on $T$, it can be shown that $v_{\text {jack }}$ consistently estimates $\mathbb{V}\left(T_{n}\right)$. For example, if $T$ is a smooth function of the sample mean, then consistency holds.
3.5 Theorem. Let $\mu=\mathbb{E}\left(X_{1}\right)$ and $\sigma^{2}=\mathbb{V}\left(X_{1}\right)<\infty$ and suppose that $T_{n}=g\left(\bar{X}_{n}\right)$ where $g$ has a continuous, nonzero derivative at $\mu$. Then $\left(T_{n}-\right.$
$g(\mu)) / \sigma_{n} \rightsquigarrow N(0,1)$ where $\sigma_{n}^{2}=n^{-1}\left(g^{\prime}(\mu)\right)^{2} \sigma^{2}$. The jackknife is consistent, meaning that

$$
\begin{equation*}
\frac{v_{\text {jack }}}{\sigma_{n}^{2}} \xrightarrow{\text { a.s. }} 1 . \tag{3.6}
\end{equation*}
$$

3.7 Theorem (Efron, 1982). If $T(F)=F^{-1}(p)$ is the $p^{\text {th }}$ quantile, then the jackknife variance estimate is inconsistent. For the median $(p=1 / 2)$ we have that $v_{\mathrm{jack}} / \sigma_{n}^{2} \rightsquigarrow\left(\chi_{2}^{2} / 2\right)^{2}$ where $\sigma_{n}^{2}$ is the asymptotic variance of the sample median.
3.8 Example. Let $T_{n}=\bar{X}_{n}$. It is easy to see that $\widetilde{T}_{i}=X_{i}$. Hence, $T_{\mathrm{jack}}=T_{n}$, $b=0$ and $v_{\text {jack }}=S_{n}^{2} / n$ where $S_{n}^{2}$ is the sample variance.

There is a connection between the jackknife and the influence function. Recall that the influence function is

$$
\begin{equation*}
L_{F}(x)=\lim _{\epsilon \rightarrow 0} \frac{T\left((1-\epsilon) F+\epsilon \delta_{x}\right)-T(F)}{\epsilon} . \tag{3.9}
\end{equation*}
$$

Suppose we approximate $L_{F}\left(X_{i}\right)$ by setting $F=\widehat{F}_{n}$ and $\epsilon=-1 /(n-1)$. This yields the approximation

$$
L_{F}\left(X_{i}\right) \approx \frac{T\left((1-\epsilon) \widehat{F}_{n}+\epsilon \delta_{x_{i}}\right)-T\left(\widehat{F}_{n}\right)}{\epsilon}=(n-1)\left(T_{n}-T_{(-i)}\right) \equiv \ell_{i}
$$

It follows that

$$
b=-\frac{1}{n} \sum_{i=1}^{n} \ell_{i}, \quad v_{\mathrm{jack}}=\frac{1}{n(n-1)}\left(\sum_{i} \ell_{i}^{2}-n b^{2}\right) .
$$

In other words, the jackknife is an approximate version of the nonparametric delta method.
3.10 Example. Consider estimating the skewness $T(F)=\int(x-\mu)^{3} d F(x) / \sigma^{3}$ of the nerve data. The point estimate is $T\left(\widehat{F}_{n}\right)=1.76$. The jackknife estimate of the standard error is .17 . An approximate 95 percent confidence interval for $T(F)$ is $1.76 \pm 2(.17)=(1.42,2.10)$. These exclude 0 which suggests that the data are not Normal. We can also compute the standard error using the influence function. For this functional, we have (see Exercise 1)

$$
L_{F}(x)=\frac{(x-\mu)^{3}}{\sigma^{3}}-T(F)\left(1+\frac{3}{2} \frac{\left((x-\mu)^{2}-\sigma^{2}\right)}{\sigma^{2}}\right)
$$

Then

$$
\widehat{\mathrm{se}}=\sqrt{\frac{\widehat{\tau}^{2}}{n}}=\sqrt{\frac{\sum_{i=1}^{n} \widehat{L}^{2}\left(X_{i}\right)}{n^{2}}}=.18
$$

It is reassuring to get nearly the same answer.

### 3.2 The Bootstrap

The bootstrap is a method for estimating the variance and the distribution of a statistic $T_{n}=g\left(X_{1}, \ldots, X_{n}\right)$. We can also use the bootstrap to construct confidence intervals.

Let $\mathbb{V}_{F}\left(T_{n}\right)$ denote the variance of $T_{n}$. We have added the subscript $F$ to emphasize that the variance is a function of $F$. If we knew $F$ we could, at least in principle, compute the variance. For example, if $T_{n}=n^{-1} \sum_{i=1}^{n} X_{i}$, then

$$
\mathbb{V}_{F}\left(T_{n}\right)=\frac{\sigma^{2}}{n}=\frac{\int x^{2} d F(x)-\left(\int x d F(x)\right)^{2}}{n}
$$

which is clearly a function of $F$.
With the bootstrap, we estimate $\mathbb{V}_{F}\left(T_{n}\right)$ with $\mathbb{V}_{\widehat{F}_{n}}\left(T_{n}\right)$. In other words, we use a plug-in estimator of the variance. Since, $\mathbb{V}_{\widehat{F}_{n}}\left(T_{n}\right)$ may be difficult to compute, we approximate it with a simulation estimate denoted by $v_{\text {boot }}$. Specifically, we do the following steps:

## Bootstrap Variance Estimation

1. Draw $X_{1}^{*}, \ldots, X_{n}^{*} \sim \widehat{F}_{n}$.
2. Compute $T_{n}^{*}=g\left(X_{1}^{*}, \ldots, X_{n}^{*}\right)$.
3. Repeat steps 1 and $2, B$ times to get $T_{n, 1}^{*}, \ldots, T_{n, B}^{*}$.
4. Let

$$
\begin{equation*}
v_{\text {boot }}=\frac{1}{B} \sum_{b=1}^{B}\left(T_{n, b}^{*}-\frac{1}{B} \sum_{r=1}^{B} T_{n, r}^{*}\right)^{2} . \tag{3.11}
\end{equation*}
$$

By the law of large numbers, $v_{\text {boot }} \xrightarrow{\text { a.s. }} \mathbb{V}_{\widehat{F}_{n}}\left(T_{n}\right)$ as $B \rightarrow \infty$. The estimated standard error of $T_{n}$ is $\widehat{\mathrm{se}}_{\text {boot }}=\sqrt{v_{\text {boot }}}$. The following diagram illustrates the bootstrap idea:

| Real world: | $F$ |
| ---: | :---: |
| Bootstrap world: | $\widehat{F}_{n} \Longrightarrow X_{1}, \ldots, X_{n}$ |$\quad \Longrightarrow T_{n}=g\left(X_{1}, \ldots, X_{n}\right)$

$$
\mathbb{V}_{F}\left(T_{n}\right) \overbrace{\overbrace{}^{O(1 / \sqrt{n})}}^{\mathbb{V}_{\widehat{F}_{n}}\left(T_{n}\right)} \overbrace{\overbrace{}^{O(1 / \sqrt{B})}} v_{\text {boot }} .
$$

How do we simulate from $\widehat{F}_{n}$ ? Since $\widehat{F}_{n}$ gives probability $1 / n$ to each data point, drawing $n$ points at random from $\widehat{F}_{n}$ is the same as drawing a

```
                    Bootstrap for the Median
Given data X = (X(1), ..., X(n)):
T = median(X)
Tboot = vector of length B
for(i in 1:N){
    Xstar = sample of size n from X (with replacement)
    Tboot[i] = median(Xstar)
    }
se = sqrt(variance(Tboot))
```

FIGURE 3.1. Pseudo-code for bootstrapping the median.
sample of size $n$ with replacement from the original data. Therefore step 1 can be replaced by:

1. Draw $X_{1}^{*}, \ldots, X_{n}^{*}$ with replacement from $X_{1}, \ldots, X_{n}$.
3.12 Example. Figure 3.1 shows pseudo-code for using the bootstrap to estimate the standard error of the median.

The bootstrap can be used to approximate the CDF of a statistic $T_{n}$. Let $G_{n}(t)=\mathbb{P}\left(T_{n} \leq t\right)$ be the CDF of $T_{n}$. The bootstrap approximation to $G_{n}$ is

$$
\begin{equation*}
\widehat{G}_{n}^{*}(t)=\frac{1}{B} \sum_{b=1}^{B} I\left(T_{n, b}^{*} \leq t\right) \tag{3.13}
\end{equation*}
$$

### 3.3 Parametric Bootstrap

So far, we have estimated $F$ nonparametrically. There is also a parametric bootstrap. If $F_{\theta}$ depends on a parameter $\theta$ and $\widehat{\theta}$ is an estimate of $\theta$, then we simply sample from $F_{\widehat{\theta}}$ instead of $\widehat{F}_{n}$. This is just as accurate, but much simpler than, the delta method.
3.14 Example. When applied to the nerve data, the bootstrap, based on $B=$ 1000 replications, yields a standard error for the estimated skewness of . 16 which is nearly identical to the jackknife estimate.

### 3.4 Bootstrap Confidence Intervals

There are several ways to construct bootstrap confidence intervals. They vary in ease of calculation and accuracy.

Normal Interval. The simplest is the Normal interval

$$
T_{n} \pm z_{\alpha / 2} \widehat{\text { se }}_{\text {boot }}
$$

where $\widehat{\mathrm{se}}_{\text {boot }}$ is the bootstrap estimate of the standard error. This interval is not accurate unless the distribution of $T_{n}$ is close to Normal.

Pivotal Intervals. Let $\theta=T(F)$ and $\widehat{\theta}_{n}=T\left(\widehat{F}_{n}\right)$ and define the pivot $R_{n}=\widehat{\theta}_{n}-\theta$. Let $H(r)$ denote the CDF of the pivot:

$$
H(r)=\mathbb{P}_{F}\left(R_{n} \leq r\right)
$$

Let $C_{n}^{\star}=(a, b)$ where

$$
a=\widehat{\theta}_{n}-H^{-1}\left(1-\frac{\alpha}{2}\right) \quad \text { and } \quad b=\widehat{\theta}_{n}-H^{-1}\left(\frac{\alpha}{2}\right) .
$$

It follows that

$$
\begin{aligned}
\mathbb{P}(a \leq \theta \leq b) & =\mathbb{P}\left(\widehat{\theta}_{n}-b \leq R_{n} \leq \widehat{\theta}_{n}-a\right) \\
& =H\left(\widehat{\theta}_{n}-a\right)-H\left(\widehat{\theta}_{n}-b\right) \\
& =H\left(H^{-1}\left(1-\frac{\alpha}{2}\right)\right)-H\left(H^{-1}\left(\frac{\alpha}{2}\right)\right) \\
& =1-\frac{\alpha}{2}-\frac{\alpha}{2}=1-\alpha .
\end{aligned}
$$

Hence, $C_{n}^{\star}$ is an exact $1-\alpha$ confidence interval for $\theta$. Unfortunately, $a$ and $b$ depend on the unknown distribution $H$ but we can form a bootstrap estimate of $H$ :

$$
\widehat{H}(r)=\frac{1}{B} \sum_{b=1}^{B} I\left(R_{n, b}^{*} \leq r\right)
$$

where $R_{n, b}^{*}=\widehat{\theta}_{n, b}^{*}-\widehat{\theta}_{n}$. Let $r_{\beta}^{*}$ denote the $\beta$ sample quantile of $\left(R_{n, 1}^{*}, \ldots, R_{n, B}^{*}\right)$ and let $\theta_{\beta}^{*}$ denote the $\beta$ sample quantile of $\left(\theta_{n, 1}^{*}, \ldots, \theta_{n, B}^{*}\right)$. Note that $r_{\beta}^{*}=$ $\theta_{\beta}^{*}-\widehat{\theta}_{n}$. It follows that an approximate $1-\alpha$ confidence interval is $C_{n}=(\widehat{a}, \widehat{b})$ where

$$
\begin{aligned}
& \widehat{a}=\widehat{\theta}_{n}-\widehat{H}^{-1}\left(1-\frac{\alpha}{2}\right)=\widehat{\theta}_{n}-r_{1-\alpha / 2}^{*}=2 \widehat{\theta}_{n}-\theta_{1-\alpha / 2}^{*} \\
& \widehat{b}=\widehat{\theta}_{n}-\widehat{H}^{-1}\left(\frac{\alpha}{2}\right)=\widehat{\theta}_{n}-r_{\alpha / 2}^{*}=2 \widehat{\theta}_{n}-\theta_{\alpha / 2}^{*} .
\end{aligned}
$$

In summary:
The $1-\alpha$ bootstrap pivotal confidence interval is

$$
\begin{equation*}
C_{n}=\left(2 \widehat{\theta}_{n}-\widehat{\theta}_{((1-\alpha / 2) B)}^{*}, 2 \widehat{\theta}_{n}-\widehat{\theta}_{((\alpha / 2) B)}^{*}\right) . \tag{3.15}
\end{equation*}
$$

Typically, this is a pointwise, asymptotic confidence interval.

The next result follows from Theorem 3.21.
3.16 Theorem. If $T(F)$ is Hadamard differentiable and $C_{n}$ is given in (3.15) then $\mathbb{P}_{F}\left(T(F) \in C_{n}\right) \rightarrow 1-\alpha$.

Studentized Pivotal Interval. There is a different version of the pivotal interval which has some advantages. Let

$$
Z_{n}=\frac{T_{n}-\theta}{\widehat{\mathrm{se}}_{\text {boot }}}
$$

and

$$
Z_{n, b}^{*}=\frac{T_{n, b}^{*}-T_{n}}{\widehat{\mathrm{se}_{b}^{*}}}
$$

where $\widehat{\mathrm{se}}_{b}^{*}$ is an estimate of the standard error of $T_{n, b}^{*}$ not $T_{n}$. Now we reason as in the pivotal interval. The sample quantiles of the bootstrap quantities $Z_{n, 1}^{*}, \ldots, Z_{n, B}^{*}$ should approximate the true quantiles of the distribution of $Z_{n}$. Let $z_{\alpha}^{*}$ denote the $\alpha$ sample quantile of $Z_{n, 1}^{*}, \ldots, Z_{n, B}^{*}$, then $\mathbb{P}\left(Z_{n} \leq z_{\alpha}^{*}\right) \approx \alpha$. Let

$$
C_{n}=\left(T_{n}-z_{1-\alpha / 2}^{*} \widehat{\operatorname{se}}_{\text {boot }}, T_{n}-z_{\alpha / 2}^{*} \widehat{\operatorname{se}}_{\text {boot }}\right) .
$$

Then,

$$
\begin{aligned}
\mathbb{P}\left(\theta \in C_{n}\right) & =\mathbb{P}\left(T_{n}-z_{1-\alpha / 2}^{*} \widehat{\mathrm{~s}}_{\mathrm{boot}} \leq \theta \leq T_{n}-z_{\alpha / 2}^{*} \widehat{\mathrm{se}}_{\mathrm{boot}}\right) \\
& =\mathbb{P}\left(z_{\alpha / 2}^{*} \leq \frac{T_{n}-\theta}{\mathrm{se}_{\mathrm{boot}}} \leq z_{1-\alpha / 2}^{*}\right) \\
& =\mathbb{P}\left(z_{\alpha / 2}^{*} \leq Z_{n} \leq z_{1-\alpha / 2}^{*}\right) \\
& \approx 1-\alpha
\end{aligned}
$$

This interval has higher accuracy then all the intervals discussed so far (see Section 3.5) but there is a catch: you need to compute $\widehat{\mathrm{se}}_{b}^{*}$ for each bootstrap sample. This may require doing a second bootstrap within each bootstrap.

The $1-\alpha$ bootstrap studentized pivotal interval is

$$
\left(T_{n}-z_{1-\alpha / 2}^{*} \widehat{\mathrm{se}}_{\text {boot }}, T_{n}-z_{\alpha / 2}^{*} \widehat{\mathrm{se}}_{\text {boot }}\right)
$$

where $z_{\beta}^{*}$ is the $\beta$ quantile of $Z_{n, 1}^{*}, \ldots, Z_{n, B}^{*}$ and

$$
Z_{n, b}^{*}=\frac{T_{n, b}^{*}-T_{n}}{\widehat{\mathrm{se}}_{b}^{*}}
$$

Percentile Intervals. The bootstrap percentile interval is defined by

$$
C_{n}=\left(T_{(B \alpha / 2)}^{*}, T_{(B(1-\alpha / 2))}^{*}\right),
$$

that is, just use the $\alpha / 2$ and $1-\alpha / 2$ quantiles of the bootstrap sample. The justification for this interval is as follows. Suppose there exists a monotone transformation $U=m(T)$ such that $U \sim N\left(\phi, c^{2}\right)$ where $\phi=m(\theta)$. We do not suppose we know the transformation, only that one exists. Let $U_{b}^{*}=m\left(T_{b}^{*}\right)$. Note that $U_{(B \alpha / 2)}^{*}=m\left(T_{(B \alpha / 2)}^{*}\right)$ since a monotone transformation preserves quantiles. Since, $U \sim N\left(\phi, c^{2}\right)$, the $\alpha / 2$ quantile of $U$ is $\phi-z_{\alpha / 2} c$. Hence, $U_{(B \alpha / 2)}^{*}=\phi-z_{\alpha / 2} c \approx U-z_{\alpha / 2} c$ and $U_{(B(1-\alpha / 2))}^{*} \approx U+z_{\alpha / 2} c$. Therefore,

$$
\begin{aligned}
\mathbb{P}\left(T_{B \alpha / 2}^{*} \leq \theta \leq T_{B(1-\alpha / 2)}^{*}\right) & =\mathbb{P}\left(m\left(T_{B \alpha / 2}^{*}\right) \leq m(\theta) \leq m\left(T_{B(1-\alpha / 2)}^{*}\right)\right) \\
& =\mathbb{P}\left(U_{B \alpha / 2}^{*} \leq \phi \leq U_{B(1-\alpha / 2)}^{*}\right) \\
& \approx \mathbb{P}\left(U-c z_{\alpha / 2} \leq \phi \leq U+c z_{\alpha / 2}\right) \\
& =\mathbb{P}\left(-z_{\alpha / 2} \leq \frac{U-\phi}{c} \leq z_{\alpha / 2}\right) \\
& =1-\alpha .
\end{aligned}
$$

Amazingly, we never need to know $m$. Unfortunately, an exact normalizing transformation will rarely exist but there may exist approximate normalizing transformations. This has led to an industry of adjusted percentile methods, the most popular being the $\mathrm{BC}_{a}$ interval (bias-corrected and accelerated). We will not consider these intervals here.
3.17 Example. For estimating the skewness of the nerve data, here are the various confidence intervals.

| Method | $95 \%$ Interval |
| :--- | :--- |
| Normal | $(1.44,2.09)$ |
| percentile | $(1.42,2.03)$ |
| pivotal | $(1.48,2.11)$ |
| studentized | $(1.45,2.28)$ |

The studentized interval requires some explanation. For each bootstrap replication we compute $\widehat{\theta}^{*}$ and we also need the standard error $\widehat{\mathrm{se}}^{*}$ of $\widehat{\theta}^{*}$. We could do a bootstrap within the bootstrap (called a double bootstrap) but this is computationally expensive. Instead, we computed $\widehat{s e}^{*}$ using the nonparametric delta method applied to the bootstrap sample as described in Example 3.10.

### 3.5 Some Theory

Under certain conditions, $\widehat{G}_{n}^{*}$ is a consistent estimate of $G_{n}(t)=\mathbb{P}\left(T_{n} \leq t\right)$. To make this precise, let $\mathbb{P}_{\widehat{F}_{n}}(\cdot)$ denote probability statements made from $\widehat{F}_{n}$, treating the original data $X_{1}, \ldots, X_{n}$ as fixed. Assume that $T_{n}=T\left(\widehat{F}_{n}\right)$ is some functional of $\widehat{F}_{n}$. Then,

$$
\begin{equation*}
\widehat{G}_{n}^{*}(t)=\mathbb{P}_{\widehat{F}_{n}}\left(T\left(\widehat{F}_{n}^{*}\right) \leq t\right)=\mathbb{P}_{\widehat{F}_{n}}\left(\sqrt{n}\left(T\left(\widehat{F}_{n}^{*}\right)-T(F)\right) \leq u\right) \tag{3.18}
\end{equation*}
$$

where $u=\sqrt{n}(t-T(F))$. Consistency of the bootstrap can now be expressed as follows.
3.19 Theorem. Suppose that $\mathbb{E}\left(X_{1}^{2}\right)<\infty$. Let $T_{n}=g\left(\bar{X}_{n}\right)$ where $g$ is continuously differentiable at $\mu=\mathbb{E}\left(X_{1}\right)$ and that $g^{\prime}(\mu) \neq 0$. Then,

$$
\begin{equation*}
\sup _{u} \mid \mathbb{P}_{\widehat{F}_{n}}\left(\sqrt{n}\left(T\left(\widehat{F}_{n}^{*}\right)-T\left(\widehat{F}_{n}\right)\right) \leq u\right)-\mathbb{P}_{F}\left(\sqrt{n}\left(T\left(\widehat{F}_{n}\right)-T(F) \leq u\right) \mid \xrightarrow{\text { a.s. }} 0 .\right. \tag{3.20}
\end{equation*}
$$

3.21 Theorem. Suppose that $T(F)$ is Hadamard differentiable with respect to $d(F, G)=\sup _{x}|F(x)-G(x)|$ and that $0<\int L_{F}^{2}(x) d F(x)<\infty$. Then,

$$
\begin{equation*}
\sup _{u}\left|\mathbb{P}_{\widehat{F}_{n}}\left(\sqrt{n}\left(T\left(\widehat{F}_{n}^{*}\right)-T\left(\widehat{F}_{n}\right)\right) \leq u\right)-\mathbb{P}_{F}\left(\sqrt{n}\left(T\left(\widehat{F}_{n}\right)-T(F)\right) \leq u\right)\right| \xrightarrow{\mathrm{P}} 0 . \tag{3.22}
\end{equation*}
$$

Look closely at Theorems 3.19 and 3.21 . It is because of results like these that the bootstrap "works." In particular, the validity of bootstrap confidence intervals depends on these theorems. See, for example, Theorem 3.16. There is a tendency to treat the bootstrap as a panacea for all problems. But the bootstrap requires regularity conditions to yield valid answers. It should not be applied blindly.

It can also be shown that the bootstrap variance estimate is consistent with some conditions on $T$. Generally, the conditions for consistency of the bootstrap are weaker than those for the jackknife. For example, the bootstrap estimate of the variance of the median is consistent, but the jackknife estimate of the variance of the median is not consistent (Theorem 3.7).

Let us now compare the accuracy of the different confidence interval methods. Consider a $1-\alpha$ one-sided interval $\left[\widehat{\theta}_{\alpha}, \infty\right)$. We would like $\mathbb{P}\left(\theta \leq \widehat{\theta}_{\alpha}\right)=\alpha$ but usually this holds only approximately. If $\mathbb{P}\left(\theta \leq \widehat{\theta}_{\alpha}\right)=\alpha+O\left(n^{-1 / 2}\right)$ then we say that the interval is first-order accurate. If $\mathbb{P}\left(\theta \leq \widehat{\theta}_{\alpha}\right)=\alpha+O\left(n^{-1}\right)$ then we say that the interval is second-order accurate. Here is the comparison:

| Method | Accuracy |
| :--- | :--- |
| Normal interval | first-order accurate |
| basic pivotal interval | first-order accurate |
| percentile interval | first-order accurate |
| studentized pivotal interval | second-order accurate |
| adjusted percentile interval | second-order accurate |

Here is an explanation of why the studentized interval is more accurate. See Davison and Hinkley (1997), and Hall (1992a), for more details. Let $Z_{n}=$ $\sqrt{n}\left(T_{n}-\theta\right) / \sigma$ be a standardized quantity that converges to a standard Normal. Thus $\mathbb{P}_{F}\left(Z_{n} \leq z\right) \rightarrow \Phi(z)$. In fact,

$$
\begin{equation*}
\mathbb{P}_{F}\left(Z_{n} \leq z\right)=\Phi(z)+\frac{1}{\sqrt{n}} a(z) \phi(z)+O\left(\frac{1}{n}\right) \tag{3.23}
\end{equation*}
$$

for some polynomial $a$ involving things like skewness. The bootstrap version satisfies

$$
\begin{equation*}
\mathbb{P}_{\widehat{F}}\left(Z_{n}^{*} \leq z\right)=\Phi(z)+\frac{1}{\sqrt{n}} \widehat{a}(z) \phi(z)+O_{P}\left(\frac{1}{n}\right) \tag{3.24}
\end{equation*}
$$

where $\widehat{a}(z)-a(z)=O_{P}\left(n^{-1 / 2}\right)$. Subtracting, we get

$$
\begin{equation*}
\mathbb{P}_{F}\left(Z_{n} \leq z\right)-\mathbb{P}_{\widehat{F}}\left(Z_{n}^{*} \leq z\right)=O_{P}\left(\frac{1}{n}\right) \tag{3.25}
\end{equation*}
$$

Now suppose we work with the nonstudentized quantity $V_{n}=\sqrt{n}\left(T_{n}-\theta\right)$. Then,

$$
\begin{aligned}
\mathbb{P}_{F}\left(V_{n} \leq z\right) & =\mathbb{P}_{F}\left(\frac{V_{n}}{\sigma} \leq \frac{z}{\sigma}\right) \\
& =\Phi\left(\frac{z}{\sigma}\right)+\frac{1}{\sqrt{n}} b\left(\frac{z}{\sigma}\right) \phi\left(\frac{z}{\sigma}\right)+O\left(\frac{1}{n}\right)
\end{aligned}
$$

for some polynomial $b$. For the bootstrap we have

$$
\begin{aligned}
\mathbb{P}_{\widehat{F}}\left(V_{n}^{*} \leq z\right) & =\mathbb{P}_{\widehat{F}}\left(\frac{V_{n}}{\widehat{\sigma}} \leq \frac{z}{\widehat{\sigma}}\right) \\
& =\Phi\left(\frac{z}{\hat{\sigma}}\right)+\frac{1}{\sqrt{n}} \widehat{b}\left(\frac{z}{\hat{\sigma}}\right) \phi\left(\frac{z}{\hat{\sigma}}\right)+O_{P}\left(\frac{1}{n}\right)
\end{aligned}
$$

where $\widehat{\sigma}=\sigma+O_{P}\left(n^{-1 / 2}\right)$. Subtracting, we get

$$
\begin{equation*}
\mathbb{P}_{F}\left(V_{n} \leq z\right)-\mathbb{P}_{\widehat{F}}\left(V_{n}^{*} \leq z\right)=O_{P}\left(\frac{1}{\sqrt{n}}\right) \tag{3.26}
\end{equation*}
$$

which is less accurate than (3.25).

### 3.6 Bibliographic Remarks

The jackknife was invented by Quenouille (1949) and Tukey (1958). The bootstrap was invented by Efron (1979). There are several books on these topics including Efron and Tibshirani (1993), Davison and Hinkley (1997), Hall (1992a) and Shao and Tu (1995). Also, see Section 3.6 of van der Vaart and Wellner (1996).

### 3.7 Appendix

The book by Shao and Tu (1995) gives an explanation of the techniques for proving consistency of the jackknife and the bootstrap. Following Section 3.1 of their text, let us look at two ways of showing that the bootstrap is consistent for the case $T_{n}=\bar{X}_{n}=n^{-1} \sum_{i=1}^{n} X_{i}$. Let $X_{1}, \ldots, X_{n} \sim F$ and let $T_{n}=\sqrt{n}\left(\bar{X}_{n}-\mu\right)$ where $\mu=\mathbb{E}\left(X_{1}\right)$. Let $H_{n}(t)=\mathbb{P}_{F}\left(T_{n} \leq t\right)$ and let $\widehat{H}_{n}(t)=$ $\mathbb{P}_{\widehat{F}_{n}}\left(T_{n}^{*} \leq t\right)$ be the bootstrap estimate of $H_{n}$ where $T_{n}^{*}=\sqrt{n}\left(\bar{X}_{n}^{*}-\bar{X}_{n}\right)$ and $X_{1}^{*}, \ldots, X_{n}^{*} \sim \widehat{F}_{n}$. Our goal is to show that $\sup _{x}\left|H_{n}(x)-\widehat{H}_{n}(x)\right| \xrightarrow{\text { a.s. }} 0$.

The first method was used by Bickel and Freedman (1981) and is based on Mallows' metric. If $X$ and $Y$ are random variables with distributions $F$ and $G$, Mallows' metric is defined by $d_{r}(F, G)=d_{r}(X, Y)=\inf \left(\mathbb{E}|X-Y|^{r}\right)^{1 / r}$ where the infimum is over all joint distributions with marginals $F$ and $G$. Here are some facts about $d_{r}$. Let $X_{n} \sim F_{n}$ and $X \sim F$. Then, $d_{r}\left(F_{n}, F\right) \rightarrow 0$ if and only if $X_{n} \rightsquigarrow X$ and $\int|x|^{r} d F_{n}(x) \rightarrow \int|x|^{r} d F(x)$. If $\mathbb{E}\left(\left|X_{1}\right|^{r}\right)<\infty$ then $d_{r}\left(\widehat{F}_{n}, F\right) \xrightarrow{\text { a.s. }} 0$. For any constant $a, d_{r}(a X, a Y)=|a| d_{r}(X, Y)$. If $\mathbb{E}\left(X^{2}\right)<\infty$
and $\mathbb{E}\left(Y^{2}\right)<\infty$ then $d_{2}(X, Y)^{2}=\left(d_{2}(X-\mathbb{E}(X), Y-\mathbb{E}(Y))\right)^{2}+|\mathbb{E}(X-Y)|^{2}$. If $\mathbb{E}\left(X_{j}\right)=\mathbb{E}\left(Y_{j}\right)$ and $\mathbb{E}\left(\left|X_{j}\right|^{r}\right)<\infty, \mathbb{E}\left(\left|Y_{j}\right|^{r}\right)<\infty$ then

$$
\left(d_{2}\left(\sum_{j=1}^{m} X_{j}, \sum_{j=1}^{m} Y_{j}\right)\right)^{2} \leq \sum_{j=1}^{m} d_{2}\left(X_{j}, Y_{j}\right)^{2} .
$$

Using the properties of $d_{r}$ we have

$$
\begin{aligned}
& d_{2}\left(\widehat{H}_{n}, H_{n}\right)=d_{2}\left(\sqrt{n}\left(\bar{X}_{n}^{*}-\bar{X}_{n}\right), \sqrt{n}\left(\bar{X}_{n}-\mu\right)\right) \\
&=\frac{1}{\sqrt{n}} d_{2}\left(\sum_{i=1}^{n}\left(X_{i}^{*}-\bar{X}_{n}\right), \sum_{i=1}^{n}\left(X_{i}-\mu\right)\right) \\
& \leq \sqrt{\frac{1}{n} \sum_{i=1}^{n} d_{2}\left(X_{i}^{*}-\bar{X}_{n}, X_{i}-\mu\right)^{2}} \\
&=d_{2}\left(X_{1}^{*}-X_{1}, X_{1}-\mu\right) \\
&=\sqrt{d_{2}\left(X_{1}^{*}, X_{1}\right)^{2}-\left(\mu-\mathbb{E}_{*} X_{1}^{*}\right)^{2}} \\
&=\sqrt{d_{2}\left(\widehat{F}_{n}, F\right)^{2}-\left(\mu-\bar{X}_{n}\right)^{2}} \\
& \xrightarrow{\text { a.s. }} 0
\end{aligned}
$$

since $d_{2}\left(\widehat{F}_{n}, F\right) \xrightarrow{\text { a.s. }} 0$ and $\bar{X}_{n} \xrightarrow{\text { a.s. }} \mu$. Hence, $\sup _{x}\left|H_{n}(x)-\widehat{H}_{n}(x)\right| \xrightarrow{\text { a.s. }} 0$.
The second method, due to Singh (1981), uses the Berry-Esséen bound (1.28) which we now review. Let $X_{1}, \ldots, X_{n}$ be IID with finite mean $\mu=$ $\mathbb{E}\left(X_{1}\right)$, variance $\sigma^{2}=\mathbb{V}\left(X_{1}\right)$ and third moment, $\mathbb{E}\left|X_{1}\right|^{3}<\infty$. Let $Z_{n}=$ $\sqrt{n}\left(\bar{X}_{n}-\mu\right) / \sigma$. Then

$$
\begin{equation*}
\sup _{z}\left|\mathbb{P}\left(Z_{n} \leq z\right)-\Phi(z)\right| \leq \frac{33}{4} \frac{\mathbb{E}\left|X_{1}-\mu\right|^{3}}{\sqrt{n} \sigma^{3}} \tag{3.27}
\end{equation*}
$$

Let $Z_{n}^{*}=\left(\bar{X}_{n}^{*}-\bar{X}_{n}\right) / \widehat{\sigma}$ where $\widehat{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}_{n}\right)^{2}$. Replacing $F$ with $\widehat{F}_{n}$ and $\bar{X}_{n}$ with $\bar{X}_{n}^{*}$ we get

$$
\begin{equation*}
\sup _{z}\left|\mathbb{P}_{\widehat{F}_{n}}\left(Z_{n}^{*} \leq z\right)-\Phi(z)\right| \leq \frac{33}{4} \frac{\sum_{i=1}^{n}\left|X_{i}-\bar{X}_{n}\right|^{3}}{n^{3 / 2} \widehat{\sigma}^{3}} \tag{3.28}
\end{equation*}
$$

Let $d(F, G)=\sup _{x}|F(x)-G(x)|$ and define $\Phi_{a}(x)=\Phi(x / a)$. Then

$$
\begin{aligned}
\sup _{z}\left|\mathbb{P}_{\widehat{F}_{n}}\left(Z_{n}^{*} \leq z\right)-\Phi(z)\right| & =\sup _{z}\left|\mathbb{P}_{\widehat{F}_{n}}\left(\sqrt{n}\left(\bar{X}_{n}^{*}-\bar{X}_{n}\right) \leq z \widehat{\sigma}\right)-\Phi\left(\frac{z \widehat{\sigma}}{\widehat{\sigma}}\right)\right| \\
& =\sup _{t}\left|\mathbb{P}_{\widehat{F}_{n}}\left(\sqrt{n}\left(\bar{X}_{n}^{*}-\bar{X}_{n}\right) \leq t\right)-\Phi_{\widehat{\sigma}}(t)\right| \\
& =d\left(\widehat{H}_{n}, \Phi_{\widehat{\sigma}}\right) .
\end{aligned}
$$

By the triangle inequality

$$
\begin{equation*}
d\left(\widehat{H}_{n}, H_{n}\right) \leq d\left(\widehat{H}_{n}, \Phi_{\widehat{\sigma}}\right)+d\left(\Phi_{\widehat{\sigma}}, \Phi_{\sigma}\right)+d\left(\Phi_{\sigma}, H_{n}\right) \tag{3.29}
\end{equation*}
$$

The third term in (3.29) goes to 0 by the central limit theorem. For the second term, $d\left(\Phi_{\widehat{\sigma}}, \Phi_{\sigma}\right) \xrightarrow{\text { a.s. }} 0$ since $\widehat{\sigma}^{2} \xrightarrow{\text { a.s. }} \sigma^{2}=\mathbb{V}\left(X_{1}\right)$. The first term is bounded by the right-hand side of (3.28). Since $\mathbb{E}\left(X_{1}^{2}\right)<\infty$, this goes to 0 by the following result: if $\mathbb{E}\left|X_{1}\right|^{\delta}<\infty$ for some $0<\delta<1$ then $n^{-1 / \delta} \sum_{i=1}^{n}\left|X_{i}\right| \xrightarrow{\text { a.s. }} 0$. In conclusion, $d\left(\widehat{H}_{n}, H_{n}\right) \xrightarrow{\text { a.s. }} 0$.

### 3.8 Exercises

1. Let $T(F)=\int(x-\mu)^{3} d F(x) / \sigma^{3}$ be the skewness. Find the influence function.
2. The following data were used to illustrate the bootstrap by Bradley Efron, the inventor of the bootstrap. The data are LSAT scores (for entrance to law school) and GPA.

| LSAT | 576 | 635 | 558 | 578 | 666 | 580 | 555 | 661 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  | 651 | 605 | 653 | 575 | 545 | 572 | 594 |  |
|  |  |  |  |  |  |  |  |  |
| GPA | 3.39 | 3.30 | 2.81 | 3.03 | 3.44 | 3.07 | 3.00 | 3.43 |
|  | 3.36 | 3.13 | 3.12 | 2.74 | 2.76 | 2.88 | 3.96 |  |

Each data point is of the form $X_{i}=\left(Y_{i}, Z_{i}\right)$ where $Y_{i}=\operatorname{LSAT}_{i}$ and $Z_{i}=$ $\mathrm{GPA}_{i}$. Find the plug-in estimate of the correlation coefficient. Estimate the standard error using (i) the influence function, (ii) the jackknife and (iii) the bootstrap. Next compute a 95 percent studentized pivotal bootstrap confidence interval. You will need to compute the standard error of $T^{*}$ for every bootstrap sample.
3. Let $T_{n}=\bar{X}_{n}^{2}, \mu=\mathbb{E}\left(X_{1}\right), \alpha_{k}=\int|x-\mu|^{k} d F(x)$ and $\widehat{\alpha}_{k}=n^{-1} \sum_{i=1}^{n} \mid X_{i}-$ $\left.\bar{X}_{n}\right|^{k}$. Show that

$$
v_{\text {boot }}=\frac{4 \bar{X}_{n}^{2} \widehat{\alpha}_{2}}{n}+\frac{4 \bar{X}_{n} \widehat{\alpha}_{3}}{n^{2}}+\frac{\widehat{\alpha}_{4}}{n^{3}} .
$$

4. Prove Theorem 3.16.
5. Repeat the calculations in Example 3.17 but use a parametric bootstrap. Assume that the data are log-Normal. That is, assume that $Y \sim$ $N\left(\mu, \sigma^{2}\right)$ where $Y=\log X$. You will draw bootstrap samples $Y_{1}^{*}, \ldots, Y_{n}^{*}$ from $N\left(\widehat{\mu}, \widehat{\sigma}^{2}\right)$. Then set $X_{i}^{*}=e^{Y_{i}^{*}}$.
6. (Computer experiment.) Conduct a simulation to compare the four bootstrap confidence interval methods. Let $n=50$ and let $T(F)=\int(x-$ $\mu)^{3} d F(x) / \sigma^{3}$ be the skewness. Draw $Y_{1}, \ldots, Y_{n} \sim N(0,1)$ and set $X_{i}=$ $e^{Y_{i}}, i=1, \ldots, n$. Construct the four types of bootstrap 95 percent intervals for $T(F)$ from the data $X_{1}, \ldots, X_{n}$. Repeat this whole thing many times and estimate the true coverage of the four intervals.
7. Let

$$
X_{1}, \ldots, X_{n} \sim t_{3}
$$

where $n=25$. Let $\theta=T(F)=\left(q .75-q_{.25}\right) / 1.34$ where $q_{p}$ denotes the $p^{t h}$ quantile. Do a simulation to compare the coverage and length of the following confidence intervals for $\theta$ : (i) Normal interval with standard error from the jackknife, (ii) Normal interval with standard error from the bootstrap, (iii) bootstrap percentile interval.

Remark: The jackknife does not give a consistent estimator of the variance of a quantile.
8. Let $X_{1}, \ldots, X_{n}$ be distinct observations (no ties). Show that there are

$$
\binom{2 n-1}{n}
$$

distinct bootstrap samples.
Hint: Imagine putting $n$ balls into $n$ buckets.
9. Let $X_{1}, \ldots, X_{n}$ be distinct observations (no ties). Let $X_{1}^{*}, \ldots, X_{n}^{*}$ denote a bootstrap sample and let $\bar{X}_{n}^{*}=n^{-1} \sum_{i=1}^{n} X_{i}^{*}$. Find: $\mathbb{E}\left(\bar{X}_{n}^{*} \mid X_{1}, \ldots, X_{n}\right)$, $\mathbb{V}\left(\bar{X}_{n}^{*} \mid X_{1}, \ldots, X_{n}\right), \mathbb{E}\left(\bar{X}_{n}^{*}\right)$ and $\mathbb{V}\left(\bar{X}_{n}^{*}\right)$.
10. (Computer experiment.) Let $X_{1}, \ldots, X_{n} \sim \operatorname{Normal}(\mu, 1)$. Let $\theta=e^{\mu}$ and let $\widehat{\theta}=e^{\bar{X}}$ be the MLE. Create a data set (using $\mu=5$ ) consisting of $\mathrm{n}=100$ observations.
(a) Use the delta method to get the se and 95 percent confidence interval for $\theta$. Use the parametric bootstrap to get the se and 95 percent confidence interval for $\theta$. Use the nonparametric bootstrap to get the se and 95 percent confidence interval for $\theta$. Compare your answers.
(b) Plot a histogram of the bootstrap replications for the parametric and nonparametric bootstraps. These are estimates of the distribution of $\widehat{\theta}$. The delta method also gives an approximation to this distribution, namely, $\operatorname{Normal}\left(\widehat{\theta}, \widehat{s e}^{2}\right)$. Compare these to the true sampling distribution of $\widehat{\theta}$. Which approximation-parametric bootstrap, bootstrap or delta method-is closer to the true distribution?
11. Let $X_{1}, \ldots, X_{n} \sim \operatorname{Uniform}(0, \theta)$. The mLE is

$$
\widehat{\theta}=X_{\max }=\max \left\{X_{1}, \ldots, X_{n}\right\}
$$

Generate a data set of size 50 with $\theta=1$.
(a) Find the distribution of $\widehat{\theta}$. Compare the true distribution of $\widehat{\theta}$ to the histograms from the parametric and nonparametric bootstraps.
(b) This is a case where the nonparametric bootstrap does very poorly. In fact, we can prove that this is the case. Show that, for the parametric bootstrap $\mathbb{P}\left(\widehat{\theta}^{*}=\widehat{\theta}\right)=0$ but for the nonparametric bootstrap $\mathbb{P}\left(\widehat{\theta}^{*}=\right.$ $\widehat{\theta}) \approx .632$. Hint: Show that, $P\left(\widehat{\theta}^{*}=\widehat{\theta}\right)=1-(1-(1 / n))^{n}$. Then take the limit as $n$ gets large.
12. Suppose that 50 people are given a placebo and 50 are given a new treatment. Thirty placebo patients show improvement, while 40 treated patients show improvement. Let $\tau=p_{2}-p_{1}$ where $p_{2}$ is the probability of improving under treatment and $p_{1}$ is the probability of improving under placebo.
(a) Find the mLE of $\tau$. Find the standard error and 90 percent confidence interval using the delta method.
(b) Find the standard error and 90 percent confidence interval using the bootstrap.
13. Let $X_{1}, \ldots, X_{n} \sim F$ be IID and let $X_{1}^{*}, \ldots, X_{n}^{*}$ be a bootstrap sample from $\widehat{F}_{n}$. Let $G$ denote the marginal distribution of $X_{i}^{*}$. Note that $G(x)=\mathbb{P}\left(X_{i}^{*} \leq x\right)=\mathbb{E} \mathbb{P}\left(X_{i}^{*} \leq x \mid X_{1}, \ldots, X_{n}\right)=\mathbb{E}\left(\widehat{F}_{n}(x)\right)=F(x)$. So it appears that $X_{i}^{*}$ and $X_{i}$ have the same distribution. But in Exercise 9 we showed that $\mathbb{V}\left(\bar{X}_{n}\right) \neq \mathbb{V}\left(\bar{X}_{n}^{*}\right)$. This appears to be a contradiction. Explain.

## 4

## Smoothing: General Concepts

To estimate a curve - such as a probability density function $f$ or a regression function $r$-we must smooth the data in some way. The rest of the book is devoted to smoothing methods. In this chapter we discuss some general issues related to smoothing. There are mainly two types of problems we will study. The first is density estimation in which we have a sample $X_{1}, \ldots, X_{n}$ from a distribution $F$ with density $f$, written

$$
\begin{equation*}
X_{1}, \ldots, X_{n} \sim f \tag{4.1}
\end{equation*}
$$

and we want to estimate the probability density function $f$. The second is regression in which we have a pairs $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ where

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\epsilon_{i} \tag{4.2}
\end{equation*}
$$

where $\mathbb{E}\left(\epsilon_{i}\right)=0$, and we want to estimate the regression function $r$. We begin with some examples all of which are discussed in more detail in the chapters that follow.
4.3 Example (Density estimation). Figure 4.1 shows histograms of 1266 data points from the Sloan Digital Sky Survey (SDSS). As described on the SDSS website www. sdss.org:

Simply put, the Sloan Digital Sky Survey is the most ambitious astronomical survey project ever undertaken. The survey will map
in detail one-quarter of the entire sky, determining the positions and absolute brightnesses of more than 100 million celestial objects. It will also measure the distances to more than a million galaxies and quasars.

Each data point $X_{i}$ is a redshift, ${ }^{1}$ essentially the distance from us to a galaxy. The data lie on a "pencil beam" which means that the sample lies in a narrow tube starting at the earth and pointing out into space; see Figure 4.2. The full dataset is three-dimensional. We have extracted data along this pencil beam to make it one-dimensional. Our goal is to understand the distribution of galaxies. In particular, astronomers are interested in finding clusters of galaxies. Because the speed of light is finite, when we look at galaxies that are farther away we are seeing further back in time. By looking at how galaxies cluster as a function of redshift, we are seeing how galaxy clustering evolved over time.

We regard the redshifts $X_{1}, \ldots, X_{n}$ as a sample from a distribution $F$ with density $f$, that is,

$$
X_{1}, \ldots, X_{n} \sim f
$$

as in (4.1). One way to find galaxy clusters is to look for peaks in the density $f$. A histogram is a simple method for estimating the density $f$. The details will be given in Chapter 6 but here is a brief description. We cut the real line into intervals, or bins, and count the number of observations in each bin. The heights of the bars in the histogram are proportional to the count in each bin. The three histograms in Figure 4.1 are based on different numbers of bins. The top left histogram uses a large number of bins, the top right histogram uses fewer bins and the bottom left histogram uses fewer still. The width $h$ of the bins is a smoothing parameter. We will see that large $h$ (few bins) leads to an estimator with large bias but small variance, called oversmoothing. Small $h$ (many bins) leads to an estimator with small bias but large variance, called undersmoothing. The bottom right plot shows an estimate of the mean squared error (MSE) of the histogram estimator, which is a measure of the inaccuracy of the estimator. The estimated MSE is shown as a function of the number of bins. The top right histogram has 308 bins, which corresponds to the minimizer of the estimated MSE.

[^2]

FIGURE 4.1. Three histograms for the astronomy data. The top left histogram has too many bins. The bottom left histogram has too few bins. The top right histogram uses 308 bins, chosen by the method of cross-validation as described in Chapter 6. The lower right plot shows the estimated mean squared error (inaccuracy) versus the number of bins.

Figure 4.3 shows a more sophisticated estimator of $f$ called a kernel estimator which we will describe in Chapter 6. Again there is a smoothing parameter $h$ and the three estimators correspond to increasing $h$. The structure in the data is evident only if we smooth the data by just the right amount (top right plot).
4.4 Example (Nonparametric regression). The beginning of the universe is usually called the Big Bang. It is misleading to think of this as an event taking place in empty space. More accurately, the early universe was in a hot, dense state. Since then, the universe has expanded and cooled. The remnants of the heat from the Big Bang are still observable and are called the cosmic mi-


FIGURE 4.2. In a pencil beam sample, the positions of galaxies are recorded along a path from the Earth and outward.
crowave background (CMB) radiation. Figure 4.4 shows CMB data from the Wilkinson Microwave Anisotropy Probe (WMAP). Information on WMAP is available at http://map.gsfc.nasa.gov. The image shows the temperature at each point in the sky. This is a snapshot of the universe 379,000 years after the Big Bang. The average temperature is 2.73 degrees above absolute zero but the temperature is not constant across the sky. The fluctuations in the temperature map provide information about the early universe. Indeed, as the universe expanded, there was a tug of war between the force of expansion and contraction due to gravity. This caused acoustic waves in the hot gas (like a vibrating bowl of jello), which is why there are temperature fluctuations. The strength of the temperature fluctuations $r(x)$ at each frequency (or multipole) $x$ is called the power spectrum and this power spectrum can be used to by cosmologists to answer cosmological questions (Genovese et al. (2004)). For example, the relative abundance of different constituents of the universe (such as baryons and dark matter) correspond to peaks in the power spectrum. Through a very complicated procedure that we won't describe here, the temperature map can be reduced to a scatterplot of power versus frequency.


FIGURE 4.3. Kernel density estimators and estimated MSE for the astronomy data. Top left: oversmoothed. Top right: just right (bandwidth chosen by cross-validation). Bottom left: undersmoothed. Bottom right: estimated MSE as a function of the smoothing parameter $h$.

The first 400 data points are shown in Figure 4.5. (All 899 data points are shown in Figure 5.3.)

At this point of the analysis, the data consist of $n$ pairs $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ where $x_{i}$ is called the multipole moment and $Y_{i}$ is the estimated power spectrum of the temperature fluctuations. If $r(x)$ denotes the true power spectrum then

$$
Y_{i}=r\left(x_{i}\right)+\epsilon_{i}
$$

where $\epsilon_{i}$ is a random error with mean 0 , as in (4.2). The goal in nonparametric regression is to estimate $r$ making only minimal assumptions about $r$. The first plot in Figure 4.5 shows the data and the next three plots show a nonparametric estimator (called a local regression estimator) of $r$ with increasing smoothing parameter $h$. The structure of the data is distorted if we smooth too little or too much. The details will be explained in the next chapter.


FIGURE 4.4. The WMAP (Wilkinson Microwave Anisotropy Probe) temperature map. This is the heat left over from the Big Bang. The data show the temperature at each point in the sky. The microwave light captured in this picture is from 379,000 years after the Big Bang (13 billion years ago). The average temperature is 2.73 degrees above absolute zero. The fluctuations in the temperature map provide important information about the early universe.
4.5 Example (Nonparametric regression). Ruppert et al. (2003) describe data from a light detection and ranging (LIDAR) experiment. LIDAR is used to monitor pollutants; see Sigrist (1994). Figure 4.6 shows 221 observations. The response is the log of the ratio of light received from two lasers. The frequency of one laser is the resonance frequency of mercury while the second has a different frequency. The estimates shown here are called regressograms, which is the regression version of a histogram. We divide the horizontal axis into bins and then we take the sample average of the $Y_{i} \mathrm{~s}$ in each bin. The smoothing parameter $h$ is the width of the bins. As the binsize $h$ decreases, the estimated regression function $\widehat{r}_{n}$ goes from oversmoothing to undersmoothing.
4.6 Example (Nonparametric binary regression). This example is from Pagano and Gauvreau (1993); it also appears in Ruppert et al. (2003). The goal is to relate the presence or absence of bronchopulmonary dysplasia (BPD) with birth weight (in grams) for 223 babies. BPD is a chronic lung disease that can affect premature babies. The outcome $Y$ is binary: $Y=1$ if the baby has BPD and $Y=0$ otherwise. The covariate is $x=$ birth weight. A common parametric model for relating a binary outcome $Y$ to a covariate $x$ is logistic regression which has the form

$$
r\left(x ; \beta_{0}, \beta_{1}\right) \equiv \mathbb{P}(Y=1 \mid X=x)=\frac{e^{\beta_{0}+\beta_{1} x}}{1+e^{\beta_{0}+\beta_{1} x}}
$$



FIGURE 4.5. The first 400 data points of the CMB data. Top left: scatterplot of power versus frequency. Top right: undersmoothed. Bottom left: just right; Bottom right: oversmoothed.

The parameters $\beta_{0}$ and $\beta_{1}$ are usually estimated by maximum likelihood. The estimated function (solid line) $r\left(x ; \widehat{\beta}_{0}, \widehat{\beta}_{1}\right)$ together with the data are shown in Figure 4.7. Also shown are two nonparametric estimates. In this example, the nonparametric estimators are not terribly different from the parametric estimator. Of course, this need not always be the case.
4.7 Example (Multiple nonparametric regression). This example, from Venables and Ripley (2002), involves three covariates and one response variable. The data are 48 rock samples from a petroleum reservoir. The response is permeability (in milli-Darcies). The covariates are: the area of pores (in pixels out of 256 by 256 ), perimeter in pixels and shape (perimeter $/ \sqrt{\text { area }}$ ). The goal is to predict permeability from the three covariates. One nonparametric model is

$$
\text { permeability }=r(\text { area }, \text { perimeter }, \text { shape })+\epsilon
$$



FIGURE 4.6. The LIDAR data from Example 4.5. The estimates are regressograms, obtained by averaging the $Y_{i}$ s over bins. As we decrease the binwidth $h$, the estimator becomes less smooth.
where $r$ is a smooth function. A simpler, but less general model is the additive model

$$
\text { permeability }=r_{1}(\text { area })+r_{2}(\text { perimeter })+r_{3}(\text { shape })+\epsilon
$$

for smooth functions $r_{1}, r_{2}$, and $r_{3}$. Estimates of $r_{1}, r_{2}$ and $r_{3}$ are shown in Figure 4.8.

### 4.1 The Bias-Variance Tradeoff

Let $\widehat{f}_{n}(x)$ be an estimate of a function $f(x)$. The squared error (or $\left.L_{2}\right) \operatorname{loss}$ function is

$$
\begin{equation*}
L\left(f(x), \widehat{f}_{n}(x)\right)=\left(f(x)-\widehat{f}_{n}(x)\right)^{2} \tag{4.8}
\end{equation*}
$$



FIGURE 4.7. The BPD data from Example 4.6. The data are shown with small vertical lines. The estimates are from logistic regression (solid line), local likelihood (dashed line) and local linear regression (dotted line).

The average of this loss is called the risk or mean squared error (MSE) and is denoted by:

$$
\begin{equation*}
\operatorname{MSE}=R\left(f(x), \widehat{f}_{n}(x)\right)=\mathbb{E}\left(L\left(f(x), \widehat{f}_{n}(x)\right)\right) . \tag{4.9}
\end{equation*}
$$

The random variable in equation (4.9) is the function $\widehat{f}_{n}$ which implicitly depends on the observed data. We will use the terms risk and mSE interchangeably. A simple calculation (Exercise 2) shows that

$$
\begin{equation*}
R\left(f(x), \widehat{f_{n}}(x)\right)=\operatorname{bias}_{x}^{2}+\mathbb{V}_{x} \tag{4.10}
\end{equation*}
$$

where

$$
\operatorname{bias}_{x}=\mathbb{E}\left(\widehat{f}_{n}(x)\right)-f(x)
$$

is the bias of $\widehat{f}_{n}(x)$ and

$$
\mathbb{V}_{x}=\mathbb{V}\left(\widehat{f}_{n}(x)\right)
$$

is the variance of $\widehat{f}_{n}(x)$. In words:


FIGURE 4.8. The rock data from Example 4.7. The plots show $\widehat{r}_{1}, \widehat{r}_{2}$, and $\widehat{r}_{3}$ for the additive model $Y=\widehat{r}_{1}\left(x_{1}\right)+\widehat{r}_{2}\left(x_{2}\right)+\widehat{r}_{3}\left(x_{3}\right)+\epsilon$.

$$
\begin{equation*}
\text { RISK }=\text { MSE }=\text { BIAS }^{2}+\text { VARIANCE } . \tag{4.11}
\end{equation*}
$$

The above definitions refer to the risk at a point $x$. Now we want to summarize the risk over different values of $x$. In density estimation problems, we will use the integrated risk or integrated mean squared error defined by

$$
\begin{equation*}
R\left(f, \widehat{f}_{n}\right)=\int R\left(f(x), \widehat{f}_{n}(x)\right) d x \tag{4.12}
\end{equation*}
$$

For regression problems we can use the integrated MSE or the average mean squared error

$$
\begin{equation*}
R\left(r, \widehat{r}_{n}\right)=\frac{1}{n} \sum_{i=1}^{n} R\left(r\left(x_{i}\right), \widehat{r}_{n}\left(x_{i}\right)\right) . \tag{4.13}
\end{equation*}
$$

The average risk has the following predictive risk interpretation. Suppose the model for the data is the nonparametric regression model

$$
Y_{i}=r\left(x_{i}\right)+\epsilon_{i} .
$$

Suppose we draw a new observation $Y_{i}^{*}=r\left(x_{i}\right)+\epsilon_{i}^{*}$ at each $x_{i}$. If we predict $Y_{i}^{*}$ with $\widehat{r}_{n}\left(x_{i}\right)$ then the squared prediction error is

$$
\left(Y_{i}^{*}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}=\left(r\left(x_{i}\right)+\epsilon_{i}^{*}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2} .
$$

Define the predictive risk

$$
\text { predictive risk }=\mathbb{E}\left(\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}^{*}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}\right)
$$

Then, we have

$$
\begin{equation*}
\text { predictive risk }=R\left(r, \widehat{r}_{n}\right)+c \tag{4.14}
\end{equation*}
$$

where $c=n^{-1} \sum_{i=1}^{n} \mathbb{E}\left(\left(\epsilon_{i}^{*}\right)^{2}\right)$ is a constant. In particular, if each $\epsilon_{i}$ has variance $\sigma^{2}$ then

$$
\begin{equation*}
\text { predictive risk }=R\left(r, \widehat{r}_{n}\right)+\sigma^{2} \tag{4.15}
\end{equation*}
$$

Thus, up to a constant, the average risk and the predictive risk are the same.
The main challenge in smoothing is to determine how much smoothing to do. When the data are oversmoothed, the bias term is large and the variance is small. When the data are undersmoothed the opposite is true; see Figure 4.9. This is called the bias-variance tradeoff. Minimizing risk corresponds to balancing bias and variance.
4.16 Example. To understand the bias-variance tradeoff better, let $f$ be a PDF and consider estimating $f(0)$. Let $h$ be a small, positive number. Define

$$
p_{h} \equiv \mathbb{P}\left(-\frac{h}{2}<X<\frac{h}{2}\right)=\int_{-h / 2}^{h / 2} f(x) d x \approx h f(0)
$$

and hence

$$
f(0) \approx \frac{p_{h}}{h}
$$

Let $X$ be the number of observations in the interval $(-h / 2, h / 2)$. Then $X \sim$ $\operatorname{Binomial}\left(n, p_{h}\right)$. An estimate of $p_{h}$ is $\widehat{p}_{h}=X / n$ and hence an estimate of $f(0)$ is

$$
\begin{equation*}
\widehat{f}_{n}(0)=\frac{\widehat{p}_{h}}{h}=\frac{X}{n h} . \tag{4.17}
\end{equation*}
$$



FIGURE 4.9. The bias-variance tradeoff. The bias increases and the variance decreases with the amount of smoothing. The optimal amount of smoothing, indicated by the vertical line, minimizes the risk $=$ bias $^{2}+$ variance.

We will now show that the MSE of this estimator takes the form

$$
\begin{equation*}
\mathrm{MSE} \approx A h^{4}+\frac{B}{n h} \tag{4.18}
\end{equation*}
$$

for some constants $A$ and $B$. The first term corresponds to the squared bias and the second term corresponds to the variance.

Since $X$ is Binomial, it has mean $n p_{h}$. Now,

$$
f(x) \approx f(0)+x f^{\prime}(0)+\frac{x^{2}}{2} f^{\prime \prime}(0)
$$

So
$p_{h}=\int_{-h / 2}^{h / 2} f(x) d x \approx \int_{-h / 2}^{h / 2}\left(f(0)+x f^{\prime}(0)+\frac{x^{2}}{2} f^{\prime \prime}(0)\right) d x=h f(0)+\frac{f^{\prime \prime}(0) h^{3}}{24}$
and hence, from (4.17),

$$
\mathbb{E}\left(\widehat{f_{n}}(0)\right)=\frac{\mathbb{E}(X)}{n h}=\frac{p_{h}}{h} \approx f(0)+\frac{f^{\prime \prime}(0) h^{2}}{24}
$$

Therefore, the bias is

$$
\begin{equation*}
\text { bias }=\mathbb{E}\left(\widehat{f}_{n}(0)\right)-f(0) \approx \frac{f^{\prime \prime}(0) h^{2}}{24} \tag{4.19}
\end{equation*}
$$

To compute the variance, note that $\mathbb{V}(X)=n p_{h}\left(1-p_{h}\right)$. Therefore,

$$
\mathbb{V}\left(\widehat{f}_{n}(0)\right)=\frac{\mathbb{V}(X)}{n^{2} h^{2}}=\frac{p_{h}\left(1-p_{h}\right)}{n h^{2}} \approx \frac{p_{h}}{n h^{2}}
$$

where we have used the fact that $1-p_{h} \approx 1$ since $h$ is small. Thus,

$$
\begin{equation*}
\mathbb{V}\left(\widehat{f}_{n}(0)\right) \approx \frac{h f(0)+\frac{f^{\prime \prime}(0) h^{3}}{24}}{n h^{2}}=\frac{f(0)}{n h}+\frac{f^{\prime \prime}(0) h}{24 n} \approx \frac{f(0)}{n h} \tag{4.20}
\end{equation*}
$$

Therefore,

$$
\begin{equation*}
\mathrm{MSE}=\operatorname{bias}^{2}+\mathbb{V}\left(\widehat{f}_{n}(0)\right) \approx \frac{\left(f^{\prime \prime}(0)\right)^{2} h^{4}}{576}+\frac{f(0)}{n h} \equiv A h^{4}+\frac{B}{n h} \tag{4.21}
\end{equation*}
$$

As we smooth more, that is, as we increase $h$, the bias term increases and the variance term decreases. As we smooth less, that is, as we decrease $h$, the bias term decreases and the variance term increases. This is a typical bias-variance analysis.

### 4.2 Kernels

Throughout the rest of the book, we will often use the word "kernel." For our purposes, the word kernel refers to any smooth function $K$ such that $K(x) \geq 0$ and

$$
\begin{equation*}
\int K(x) d x=1, \quad \int x K(x) d x=0 \quad \text { and } \quad \sigma_{K}^{2} \equiv \int x^{2} K(x) d x>0 \tag{4.22}
\end{equation*}
$$

Some commonly used kernels are the following:

$$
\begin{array}{rlrl}
\text { the boxcar kernel : } & & K(x)=\frac{1}{2} I(x), \\
\text { the Gaussian kernel : } & K(x)=\frac{1}{\sqrt{2 \pi}} e^{-x^{2} / 2}, \\
\text { the Epanechnikov kernel : } & K(x)=\frac{3}{4}\left(1-x^{2}\right) I(x) \\
\text { the tricube kernel : } & K(x)=\frac{70}{81}\left(1-|x|^{3}\right)^{3} I(x)
\end{array}
$$

where

$$
I(x)= \begin{cases}1 & \text { if }|x| \leq 1 \\ 0 & \text { if }|x|>1\end{cases}
$$

These kernels are plotted in Figure 4.10.
Kernels are used for taking local averages. For example, suppose we have pairs of data $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ and we want to take the average of all the
$Y_{i}$ 's whose corresponding $x_{i}$ 's are within a distance $h$ of some point $x$. This local average is equal to

$$
\begin{equation*}
\sum_{i=1}^{n} Y_{i} \ell_{i}(x) \tag{4.23}
\end{equation*}
$$

where

$$
\begin{equation*}
\ell_{i}(x)=\frac{K\left(\frac{X_{i}-x}{h}\right)}{\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right)} \tag{4.24}
\end{equation*}
$$

and $K$ is the boxcar kernel. If we replace the boxcar kernel with another kernel, then (4.23) becomes a locally weighted average. Kernels will play an important role in many estimation methods. The smoother kernels lead to smoother estimates and are usually preferred to the boxcar kernel.


FIGURE 4.10. Examples of kernels: boxcar (top left), Gaussian (top right), Epanechnikov (bottom left), and tricube (bottom right).

### 4.3 Which Loss Function?

We can use other loss functions besides squared error. The $L_{p}$ loss

$$
\left\{\int\left|f(x)-\widehat{f}_{n}(x)\right|^{p}\right\}^{1 / p}
$$

has received some attention especially $L_{1}$ which tends to be outlier resistant and is invariant under one-to-one transformations. The results for other $L_{p}$ loss functions are not dramatically different than $L_{2}$ but $L_{p}$ is much harder to work with when $p \neq 2$. Hence, $L_{2}$ remains the popular choice. In the machine learning community there is some interest in Kullback-Leibler loss

$$
L\left(f, \widehat{f}_{n}\right)=\int f(x)\left(\log f(x) / \widehat{f}_{n}(x)\right) d x
$$

In fact, maximum likelihood estimation implicitly uses this loss function. Despite the natural role of Kullback-Leibler distance in parametric statistics, it is usually not an appropriate loss function in smoothing problems because it is exquisitely sensitive to the tails of the distribution; see Hall (1987). As a result, the tails can drive the whole estimation process.

### 4.4 Confidence Sets

Providing an estimate $\widehat{f}_{n}$ of a curve $f$ is rarely sufficient for drawing scientific conclusions. In the chapters that follow we shall also provide a confidence set for $f$. This will either take the form of a ball

$$
\mathcal{B}_{n}=\left\{f: \int\left(f(x)-\widehat{f}_{n}(x)\right)^{2} d x \leq s_{n}^{2}\right\}
$$

with some radius $s_{n}$ or a band (or envelope)

$$
\mathcal{B}_{n}=\{f: \ell(x) \leq f(x) \leq u(x), \text { for all } x\}
$$

based on a pair of functions $(\ell(x), u(x))$. In either case, we would like to have

$$
\begin{equation*}
\mathbb{P}_{f}\left(f \in \mathcal{B}_{n}\right) \geq 1-\alpha \tag{4.25}
\end{equation*}
$$

for all $f \in \mathcal{F}$ where $\mathcal{F}$ is some large class of functions. In practice, it may be difficult to find $\mathcal{B}_{n}$ so that (4.25) is true exactly. Instead, we may have to settle for some version of (4.25) that holds only approximately. The main point to keep in mind is that an estimate $\widehat{f}_{n}$ without some sort of confidence set is usually not useful.

### 4.5 The Curse of Dimensionality

A problem that occurs with smoothing methods is the curse of dimensionality, a term usually attributed to Bellman (1961). Roughly speaking, this means that estimation gets harder very quickly as the dimension of the observations increases.

There are at least two versions of this curse. The first is the computational curse of dimensionality. This refers to the fact that the computational burden of some methods can increase exponentially with dimension. Our focus here, however, is with the second version, which we call the statistical curse of dimensionality: if the data have dimension $d$, then we need a sample size $n$ that grows exponentially with $d$. Indeed, in the next few chapters we will see that the mean squared of any nonparametric estimator of a smooth (twice differentiable) curve will typically have mean squared error of the form

$$
\operatorname{MSE} \approx \frac{c}{n^{4 /(4+d)}}
$$

for some $c>0$. If we want the MSE to be equal to some small number $\delta$, we can set MSE $=\delta$ and solve for $n$. We find that

$$
n \propto\left(\frac{c}{\delta}\right)^{d / 4}
$$

which grows exponentially with dimension $d$.
The reason for this phenomenon is that smoothing involves estimating a function $f(x)$ using data points in a local neighborhood of $x$. But in a highdimensional problem, the data are very sparse, so local neighborhoods contain very few points.

Consider an example. Suppose we have $n$ data points uniformly distributed on the interval $[-1,1]$. How many data points will we find in the interval $[-0.1,0.1]$ ? The answer is: about $n / 10$ points. Now suppose we have $n$ data points on the 10 -dimensional unit cube $[0,1]^{10}=[0,1] \times \cdots \times[0,1]$. How many data points will we find in the cube $[-0.1,0.1]^{10}$ ? The answer is about

$$
n \times\left(\frac{.2}{2}\right)^{10}=\frac{n}{10,000,000,000} .
$$

Thus, $n$ has to be huge to ensure that small neighborhoods have any data in them.

The bottom line is that all the methods we will discuss can, in principle, be used in high-dimensional problems. However, even if we can overcome the
computational problems, we are still left with the statistical curse of dimensionality. You may be able to compute an estimator but it will not be accurate. In fact, if you compute a confidence interval around the estimate (as you should) it will be distressingly large. This is not a failure of the method. Rather, the confidence interval correctly indicates the intrinsic difficulty of the problem.

### 4.6 Bibliographic Remarks

There are a number of good texts on smoothing methods including Silverman (1986), Scott (1992), Simonoff (1996), Ruppert et al. (2003), and Fan and Gijbels (1996). Other references can be found at the end of Chapters 5 and 6. Hall (1987) discusses problems with Kullback-Leibler loss. For an extensive discussion on the curse of dimensionality, see Hastie et al. (2001).

### 4.7 Exercises

1. Let $X_{1}, \ldots, X_{n}$ be iid from a distribution $F$ with density $f$. The likelihood function for $f$ is

$$
\mathcal{L}_{n}(f)=\prod_{i=1}^{n} f\left(X_{i}\right)
$$

If the model is set $\mathcal{F}$ of all probability density functions, what is the maximum likelihood estimator of $f$ ?
2. Prove equation (4.10).
3. Let $X_{1}, \ldots, X_{n}$ be an IID sample from a $N(\theta, 1)$ distribution with density $f_{\theta}(x)=(2 \pi)^{-1 / 2} e^{-(x-\theta)^{2} / 2}$. Consider the density estimator $\widehat{f}(x)=$ $f_{\widehat{\theta}}(x)$ where $\widehat{\theta}=\bar{X}_{n}$ is the sample mean. Find the risk of $\widehat{f}$.
4. Recall that the Kullback-Leibler distance between two densities $f$ and $g$ is $D(f, g)=\int f(x) \log (f(x) / g(x)) d x$. Consider a one-dimensional parametric model $\left\{f_{\theta}(x): \theta \in \mathbb{R}\right\}$. Establish an approximation relationship between $L_{2}$ loss for the parameter $\theta$ and Kullback-Leibler loss. Specifically, show that $D\left(f_{\theta}, f_{\psi}\right) \approx(\theta-\psi)^{2} I(\theta) / 2$ where $\theta$ is the true value, $\psi$ is close to $\theta$ and $I(\theta)$ denotes the Fisher information.
5. What is the relationship between $L_{1}$ loss, $L_{2}$ loss and Kullback-Leibler loss?
6. Repeat the derivation leading up to equation (4.21) but take $X$ to have $d$ dimensions. Replace the small interval $[-h / 2, h / 2]$ with a small, $d-$ dimensional rectangle. Find the value of $h$ that minimizes the MSE. Find out how large $n$ needs so that MSE is equal to 0.1 .
7. Download the datasets from the examples in this chapter from the book website. Write programs to compute histograms and regressograms and try them on these datasets.

## 5

## Nonparametric Regression

In this chapter we will study nonparametric regression, also known as "learning a function" in the jargon of machine learning. We are given $n$ pairs of observations $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ as in Figures 5.1, 5.2 and 5.3. The response variable $Y$ is related to the covariate $x$ by the equations

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\epsilon_{i}, \quad \mathbb{E}\left(\epsilon_{i}\right)=0, \quad i=1, \ldots, n \tag{5.1}
\end{equation*}
$$

where $r$ is the regression function. The variable $x$ is also called a feature. We want to estimate (or "learn") the function $r$ under weak assumptions. The estimator of $r(x)$ is denoted by $\widehat{r}_{n}(x)$. We also refer to $\widehat{r}_{n}(x)$ as a smoother. At first, we will make the simplifying assumption that the variance $\mathbb{V}\left(\epsilon_{i}\right)=\sigma^{2}$ does not depend on $x$. We will relax this assumption later.
In (5.1), we are treating the covariate values $x_{i}$ as fixed. We could instead treat these as random, in which case we write the data as $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ and $r(x)$ is then interpreted as the mean of $Y$ conditional on $X=x$, that is,

$$
\begin{equation*}
r(x)=\mathbb{E}(Y \mid X=x) \tag{5.2}
\end{equation*}
$$

There is little difference in the two approaches and we shall mostly take the "fixed $x$ " approach except where noted.
5.3 Example (CMB data). Recall the CMB (cosmic microwave background radiation) data from Example 4.4.


FIGURE 5.1. CMB data. The horizontal axis is the multipole moment, essentially the frequency of fluctuations in the temperature field of the CMB. The vertical axis is the power or strength of the fluctuations at each frequency. The top plot shows the full data set. The bottom plot shows the first 400 data points. The first peak, around $x \approx 200$, is obvious. There may be a second and third peak further to the right.

Figure 5.1 shows the data. ${ }^{1}$ The first plot shows 899 data points over the whole range while the second plot shows the first 400 data points. We have noisy measurements $Y_{i}$ of $r\left(x_{i}\right)$ so the data are of the form (5.1). Our goal is to estimate $r$. The variance $\mathbb{V}\left(\epsilon_{i}\right)$ is most definitely not constant as a function of $x$. However, the second plot shows that the constant variance assumption is reasonable for the first 400 data points. It is believed that $r$ may have three peaks over the range of the data. The first peak is obvious from the second plot. The presence of a second or third peak is much less obvious; careful inferences are required to assess the significance of these peaks.

[^3]The methods we consider in this chapter are local regression methods and penalization methods. The former includes kernel regression and local polynomial regression. The latter leads to methods based on splines. In Chapters 8 and 9 we will consider a different approach based on orthogonal functions. All the estimators in this chapter are linear smoothers, a point we will discuss in Section 5.2.

Before we plunge into nonparametric regression, we first briefly review ordinary linear regression and its close relative, logistic regression. For more on linear regression, see Weisberg (1985).

### 5.1 Review of Linear and Logistic Regression

Suppose we have data $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ where $Y_{i} \in \mathbb{R}$ and $x_{i}=\left(x_{i 1}, \ldots\right.$, $\left.x_{i p}\right)^{T} \in \mathbb{R}^{p}$. The linear regression model assumes that

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\epsilon_{i} \equiv \sum_{j=1}^{p} \beta_{j} x_{i j}+\epsilon_{i}, \quad i=1, \ldots, n \tag{5.4}
\end{equation*}
$$

where $\mathbb{E}\left(\epsilon_{i}\right)=0$ and $\mathbb{V}\left(\epsilon_{i}\right)=\sigma^{2}$.
Warning! Usually, we want to include an intercept in the model so we will adopt the convention that $x_{i 1}=1$.

The design matrix $X$ is the $n \times p$ matrix defined by

$$
X=\left(\begin{array}{cccc}
x_{11} & x_{12} & \ldots & x_{1 p} \\
x_{21} & x_{22} & \ldots & x_{2 p} \\
\vdots & \vdots & \vdots & \vdots \\
x_{n 1} & x_{n 2} & \ldots & x_{n p}
\end{array}\right)
$$

The set $\mathcal{L}$ of vectors that can be obtained as linear combinations of the columns of $X$, is called the column space of $X$.

Let $Y=\left(Y_{1}, \ldots, Y_{n}\right)^{T}, \epsilon=\left(\epsilon_{1}, \ldots, \epsilon_{n}\right)^{T}$, and $\beta=\left(\beta_{1}, \ldots, \beta_{p}\right)^{T}$. We can then write (5.4) as

$$
\begin{equation*}
Y=X \beta+\epsilon \tag{5.5}
\end{equation*}
$$

The least squares estimator $\widehat{\beta}=\left(\widehat{\beta}_{1}, \ldots, \widehat{\beta}_{p}\right)^{T}$ is the vector that minimizes the residual sums of squares

$$
\operatorname{RSS}=(Y-X \beta)^{T}(Y-X \beta)=\sum_{i=1}^{n}\left(Y_{i}-\sum_{j=1}^{p} x_{i j} \beta_{j}\right)^{2}
$$

Assuming that $X^{T} X$ is invertible, the least squares estimator is

$$
\begin{equation*}
\widehat{\beta}=\left(X^{T} X\right)^{-1} X^{T} Y \tag{5.6}
\end{equation*}
$$

The estimate of $r(x)$ at $x=\left(x_{1}, \ldots, x_{p}\right)^{T}$ is thus

$$
\widehat{r}_{n}(x)=\sum_{j=1}^{p} \widehat{\beta}_{j} x_{j}=x^{T} \widehat{\beta} .
$$

It follows that the fitted values $\mathbf{r}=\left(\widehat{r}_{n}\left(x_{1}\right), \ldots, \widehat{r}_{n}\left(x_{n}\right)\right)^{T}$ can be written as

$$
\begin{equation*}
\mathbf{r}=X \widehat{\beta}=L Y \tag{5.7}
\end{equation*}
$$

where

$$
\begin{equation*}
L=X\left(X^{T} X\right)^{-1} X^{T} \tag{5.8}
\end{equation*}
$$

is called the hat matrix. The elements of the vector $\widehat{\epsilon}=Y-\mathbf{r}$ are called the residuals. The hat matrix is symmetric $L=L^{T}$ and idempotent, $L^{2}=L$. It follows that $\mathbf{r}$ is the projection of $Y$ onto the column space $\mathcal{L}$ of $X$. It can be shown that the number of parameters $p$ is related to the matrix $L$ by way of the equation

$$
\begin{equation*}
p=\operatorname{tr}(L) \tag{5.9}
\end{equation*}
$$

where $\operatorname{tr}(L)$ denotes the trace of the matrix $L$, that is, the sum of its diagonal elements. In nonparametric regression, the number of parameters will be replaced with the effective degrees of freedom which will be defined through an equation like (5.9).

Given any $x=\left(x_{1}, \ldots, x_{p}\right)^{T}$, we can write

$$
\begin{equation*}
\widehat{r}_{n}(x)=\ell(x)^{T} Y=\sum_{i=1}^{n} \ell_{i}(x) Y_{i} \tag{5.10}
\end{equation*}
$$

where

$$
\ell(x)^{T}=x^{T}\left(X^{T} X\right)^{-1} X^{T} .
$$

An unbiased estimator of $\sigma^{2}$ is

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{\sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}}{n-p}=\frac{\|\widehat{\epsilon}\|^{2}}{n-p} . \tag{5.11}
\end{equation*}
$$

Next, we construct a confidence band for $r(x)$. We want to find a pair of functions $a(x), b(x)$ such that

$$
\begin{equation*}
\mathbb{P}(a(x) \leq r(x) \leq b(x) \text { for all } x) \geq 1-\alpha . \tag{5.12}
\end{equation*}
$$

Since $\widehat{r}_{n}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i}$, it follows that

$$
\mathbb{V}\left(\widehat{r}_{n}(x)\right)=\sigma^{2} \sum_{i=1}^{n} \ell_{i}^{2}(x)=\sigma^{2}\|\ell(x)\|^{2}
$$

which suggests using bands of the form

$$
\begin{equation*}
I(x)=(a(x), b(x)) \equiv\left(\widehat{r}_{n}(x)-c \widehat{\sigma}\|\ell(x)\|, \widehat{r}_{n}(x)+c \widehat{\sigma}\|\ell(x)\|\right) \tag{5.13}
\end{equation*}
$$

for some constant $c$. The following theorem can be found in Scheffé (1959). Let $F_{p, n-p}$ denote a random variable that has an $F$ distribution with $p$ and $n-p$ degrees of freedom. Let $F_{\alpha ; p, n-p}$ denote the upper $\alpha$ quantile for this random variable, i.e., $\mathbb{P}\left(F_{p, n-p}>F_{\alpha ; p, n-p}\right)=\alpha$.
5.14 Theorem. The confidence band defined in (5.13) with $c=\sqrt{p F_{\alpha ; p, n-p}}$ satisfies (5.12).

When the $Y_{i}$ 's are not continuous, ordinary linear regression may not be appropriate. For example, suppose that $Y_{i} \in\{0,1\}$. In this case, a commonly used parametric model is the logistic regression model given by

$$
\begin{equation*}
p_{i} \equiv p_{i}(\beta)=\mathbb{P}\left(Y_{i}=1\right)=\frac{e^{\sum_{j} \beta_{j} x_{i j}}}{1+e^{\sum_{j} \beta_{j} x_{i j}}} \tag{5.15}
\end{equation*}
$$

As before, we can include an intercept by requiring that $x_{i 1}=1$ for all $i$. This model asserts that $Y_{i}$ is a Bernoulli random variable with mean $p_{i}$. The parameters $\beta=\left(\beta_{1}, \ldots, \beta_{p}\right)^{T}$ are usually estimated by maximum likelihood. Recall that if $Y \sim \operatorname{Bernoulli}(p)$ then its probability function is $\mathbb{P}(Y=y) \equiv$ $f(y)=p^{y}(1-p)^{1-y}$. Thus the likelihood function for the model (5.15) is

$$
\begin{equation*}
\mathcal{L}(\beta)=\prod_{i=1}^{n} p_{i}^{Y_{i}}\left(1-p_{i}\right)^{1-Y_{i}} \tag{5.16}
\end{equation*}
$$

The maximum likelihood estimator $\widehat{\beta}=\left(\widehat{\beta}_{1}, \ldots, \widehat{\beta}_{p}\right)^{T}$ cannot be found in closed form. However, there is an iterative algorithm called reweighted least squares which works as follows.

## Reweighted Least Squares Algorithm

Choose starting values $\widehat{\beta}=\left(\widehat{\beta}_{1}, \ldots, \widehat{\beta}_{p}\right)^{T}$ and compute $p_{i}$ using equation (5.15), for $i=1, \ldots, n$ with $\beta_{j}$ replaced by its current estimate $\widehat{\beta}_{j}$. Iterate the following steps until convergence.

1. Set

$$
Z_{i}=\log \left(\frac{p_{i}}{1-p_{i}}\right)+\frac{Y_{i}-p_{i}}{p_{i}\left(1-p_{i}\right)}, \quad i=1, \ldots, n
$$

2. Set the new estimate of $\beta$ to be

$$
\widehat{\beta}=\left(X^{T} W X\right)^{-1} X^{T} W Z
$$

where $W$ is the diagonal matrix with $(i, i)$ element equal to $p_{i}\left(1-p_{i}\right)$. This corresponds to doing a (weighted) linear regression of $Z$ on $X$.
3. Compute the $p_{i}$ s using (5.15) with the current estimate of $\widehat{\beta}$.

Logistic regression and linear regression are special cases of a class of models called generalized linear models. See McCullagh and Nelder (1999) for details.

### 5.2 Linear Smoothers

As we remarked earlier, all the nonparametric estimators in this chapter are linear smoothers. The formal definition is as follows.
5.17 Definition. An estimator $\widehat{r}_{n}$ of $r$ is a linear smoother if, for each $x$, there exists a vector $\ell(x)=\left(\ell_{1}(x), \ldots, \ell_{n}(x)\right)^{T}$ such that

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i} \tag{5.18}
\end{equation*}
$$

Define the vector of fitted values

$$
\begin{equation*}
\mathbf{r}=\left(\widehat{r}_{n}\left(x_{1}\right), \ldots, \widehat{r}_{n}\left(x_{n}\right)\right)^{T} \tag{5.19}
\end{equation*}
$$

where $Y=\left(Y_{1}, \ldots, Y_{n}\right)^{T}$. It then follows that

$$
\begin{equation*}
\mathbf{r}=L Y \tag{5.20}
\end{equation*}
$$

where $L$ is an $n \times n$ matrix whose $i^{\text {th }}$ row is $\ell\left(x_{i}\right)^{T}$; thus, $L_{i j}=\ell_{j}\left(x_{i}\right)$. The entries of the $i^{\text {th }}$ row show the weights given to each $Y_{i}$ in forming the estimate $\widehat{r}_{n}\left(x_{i}\right)$.
5.21 Definition. The matrix $L$ is called the smoothing matrix or the hat matrix. The $i^{\text {th }}$ row of $L$ is called the effective kernel for estimating $r\left(x_{i}\right)$. In analogy to (5.9), we define the effective degrees of
freedom by

$$
\begin{equation*}
\nu=\operatorname{tr}(L) \tag{5.22}
\end{equation*}
$$

Warning! The reader should not confuse linear smoothers-smoothers of the form (5.18) - with linear regression, in which one assumes that the regression function $r(x)$ is linear.
5.23 Remark. The weights in all the smoothers we will use have the property that, for all $x, \sum_{i=1}^{n} \ell_{i}(x)=1$. This implies that the smoother preserves constant curves. Thus, if $Y_{i}=c$ for all $i$, then $\widehat{r}_{n}(x)=c$.
5.24 Example (Regressogram). Suppose that $a \leq x_{i} \leq b i=1, \ldots, n$. Divide $(a, b)$ into $m$ equally spaced bins denoted by $B_{1}, B_{2}, \ldots, B_{m}$. Define $\widehat{r}_{n}(x)$ by

$$
\begin{equation*}
\widehat{r}_{n}(x)=\frac{1}{k_{j}} \sum_{i: x_{i} \in B_{j}} Y_{i}, \quad \text { for } x \in B_{j} \tag{5.25}
\end{equation*}
$$

where $k_{j}$ is the number of points in $B_{j}$. In other words, the estimate $\widehat{r}_{n}$ is a step function obtained by averaging the $Y_{i}$ s over each bin. This estimate is called the regressogram. An example is given in Figure 4.6. For $x \in B_{j}$ define $\ell_{i}(x)=1 / k_{j}$ if $x_{i} \in B_{j}$ and $\ell_{i}(x)=0$ otherwise. Thus, $\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x)$. The vector of weights $\ell(x)$ looks like this:

$$
\ell(x)^{T}=\left(0,0, \ldots, 0, \frac{1}{k_{j}}, \ldots, \frac{1}{k_{j}}, 0, \ldots, 0\right)
$$

To see what the smoothing matrix $L$ looks like, suppose that $n=9, m=3$ and $k_{1}=k_{2}=k_{3}=3$. Then,

$$
L=\frac{1}{3} \times\left(\begin{array}{ccccccccc}
1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1
\end{array}\right)
$$

In general, it is easy to see that there are $\nu=\operatorname{tr}(L)=m$ effective degrees of freedom. The binwidth $h=(b-a) / m$ controls how smooth the estimate is. and the smoothing matrix $L$ has the form
5.26 Example (Local averages). Fix $h>0$ and let $B_{x}=\left\{i:\left|x_{i}-x\right| \leq h\right\}$. Let $n_{x}$ be the number of points in $B_{x}$. For any $x$ for which $n_{x}>0$ define

$$
\widehat{r}_{n}(x)=\frac{1}{n_{x}} \sum_{i \in B_{x}} Y_{i} .
$$

This is the local average estimator of $r(x)$, a special case of the kernel estimator discussed shortly. In this case, $\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x)$ where $\ell_{i}(x)=$ $1 / n_{x}$ if $\left|x_{i}-x\right| \leq h$ and $\ell_{i}(x)=0$ otherwise. As a simple example, suppose that $n=9, x_{i}=i / 9$ and $h=1 / 9$. Then,

$$
L=\left(\begin{array}{ccccccccc}
1 / 2 & 1 / 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 / 3 & 1 / 3 & 1 / 3 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 / 3 & 1 / 3 & 1 / 3 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 / 3 & 1 / 3 & 1 / 3 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 / 3 & 1 / 3 & 1 / 3 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 / 3 & 1 / 3 & 1 / 3 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 / 3 & 1 / 3 & 1 / 3 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 / 3 & 1 / 3 & 1 / 3 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 / 2 & 1 / 2
\end{array}\right) .
$$

### 5.3 Choosing the Smoothing Parameter

The smoothers we will use will depend on some smoothing parameter $h$ and we will need some way of choosing $h$. As in Chapter 4, define the risk (mean squared error)

$$
\begin{equation*}
R(h)=\mathbb{E}\left(\frac{1}{n} \sum_{i=1}^{n}\left(\widehat{r}_{n}\left(x_{i}\right)-r\left(x_{i}\right)\right)^{2}\right) . \tag{5.27}
\end{equation*}
$$

Ideally, we would like to choose $h$ to minimize $R(h)$ but $R(h)$ depends on the unknown function $r(x)$. Instead, we will minimize an estimate $\widehat{R}(h)$ of $R(h)$. As a first guess, we might use the average residual sums of squares, also called the training error

$$
\begin{equation*}
\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2} \tag{5.28}
\end{equation*}
$$

to estimate $R(h)$. This turns out to be a poor estimate of $R(h)$ : it is biased downwards and typically leads to undersmoothing (overfitting). The reason is that we are using the data twice: to estimate the function and to estimate the risk. The function estimate is chosen to make $\sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}$ small so this will tend to underestimate the risk.

We will estimate the risk using the leave-one-out cross-validation score which is defined as follows.
5.29 Definition. The leave-one-out cross-validation score is defined by

$$
\begin{equation*}
\mathrm{CV}=\widehat{R}(h)=\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{(-i)}\left(x_{i}\right)\right)^{2} \tag{5.30}
\end{equation*}
$$

where $\widehat{r}_{(-i)}$ is the estimator obtained by omitting the $i^{\text {th }}$ pair $\left(x_{i}, Y_{i}\right)$.

As stated above, Definition 5.29 is incomplete. We have not said what we mean precisely by $\widehat{r}_{(-i)}$. We shall define

$$
\begin{equation*}
\widehat{r}_{(-i)}(x)=\sum_{j=1}^{n} Y_{j} \ell_{j,(-i)}(x) \tag{5.31}
\end{equation*}
$$

where

$$
\ell_{j,(-i)}(x)= \begin{cases}0 & \text { if } j=i  \tag{5.32}\\ \frac{\ell_{j}(x)}{\sum_{k \neq i} \ell_{k}(x)} & \text { if } j \neq i\end{cases}
$$

In other words we set the weight on $x_{i}$ to 0 and renormalize the other weights to sum to one. For all the methods in this chapter (kernel regression, local polynomials, smoothing splines) this form for $\widehat{r}_{(-i)}$ can actually be derived as a property of the method, rather than a matter of definition. But it is simpler to treat this as a definition.

The intuition for cross-validation is as follows. Note that

$$
\begin{aligned}
\mathbb{E}\left(Y_{i}-\widehat{r}_{(-i)}\left(x_{i}\right)\right)^{2} & =\mathbb{E}\left(Y_{i}-r\left(x_{i}\right)+r\left(x_{i}\right)-\widehat{r}_{(-i)}\left(x_{i}\right)\right)^{2} \\
& =\sigma^{2}+\mathbb{E}\left(r\left(x_{i}\right)-\widehat{r}_{(-i)}\left(x_{i}\right)\right)^{2} \\
& \approx \sigma^{2}+\mathbb{E}\left(r\left(x_{i}\right)-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}
\end{aligned}
$$

and hence, recalling (4.15),

$$
\begin{equation*}
\mathbb{E}(\widehat{R}) \approx R+\sigma^{2}=\text { predictive risk } \tag{5.33}
\end{equation*}
$$

Thus the cross-validation score is a nearly unbiased estimate of the risk.
It seems that it might be time consuming to evaluate $\widehat{R}(h)$ since we apparently need to recompute the estimator after dropping out each observation. Fortunately, there is a shortcut formula for computing $\widehat{R}$ for linear smoothers.
5.34 Theorem. Let $\widehat{r}_{n}$ be a linear smoother. Then the leave-one-out cross-validation score $\widehat{R}(h)$ can be written as

$$
\begin{equation*}
\widehat{R}(h)=\frac{1}{n} \sum_{i=1}^{n}\left(\frac{Y_{i}-\widehat{r}_{n}\left(x_{i}\right)}{1-L_{i i}}\right)^{2} \tag{5.35}
\end{equation*}
$$

where $L_{i i}=\ell_{i}\left(x_{i}\right)$ is the $i^{\text {th }}$ diagonal element of the smoothing matrix $L$.
The smoothing parameter $h$ can then be chosen by minimizing $\widehat{R}(h)$.
Warning! You can't assume that $\widehat{R}(h)$ always has a well-defined minimum. You should always plot $\widehat{R}(h)$ as a function of $h$.

Rather than minimizing the cross-validation score, an alternative is to use an approximation called generalized cross-validation ${ }^{2}$ in which each $L_{i i}$ in equation (5.35) is replaced with its average $n^{-1} \sum_{i=1}^{n} L_{i i}=\nu / n$ where $\nu=\operatorname{tr}(L)$ is the effective degrees of freedom. Thus, we would minimize

$$
\begin{equation*}
\operatorname{GCV}(h)=\frac{1}{n} \sum_{i=1}^{n}\left(\frac{Y_{i}-\widehat{r}_{n}\left(x_{i}\right)}{1-\nu / n}\right)^{2} \tag{5.36}
\end{equation*}
$$

Usually, the bandwidth that minimizes the generalized cross-validation score is close to the bandwidth that minimizes the cross-validation score.

Using the approximation $(1-x)^{-2} \approx 1+2 x$ we see that

$$
\begin{equation*}
\operatorname{GCV}(h) \approx \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}+\frac{2 \nu \widehat{\sigma}^{2}}{n} \equiv C_{p} \tag{5.37}
\end{equation*}
$$

where $\widehat{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}$. Equation (5.37) is known as the $C_{p}$ statistic $^{3}$ which was originally proposed by Colin Mallows as a criterion for variable selection in linear regression. More generally, many common bandwidth selection criteria can be written in the form

$$
\begin{equation*}
B(h)=\Xi(n, h) \times \frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2} \tag{5.38}
\end{equation*}
$$

for different choices of $\Xi(n, h)$. See Härdle et al. (1988) for details. Moreover, Härdle et al. (1988) prove, under appropriate conditions, the following facts about the minimizer $\widehat{h}$ of $B(h)$. Let $\widehat{h}_{0}$ minimize the loss $L(\widehat{h})=$ $n^{-1} \sum_{i=1}^{n}\left(\widehat{r}_{n}\left(x_{i}\right)-r\left(x_{i}\right)\right)^{2}$, and let $h_{0}$ minimize the risk. Then all of $\widehat{h}, \widehat{h}_{0}$, and

[^4]$h_{0}$ tend to 0 at rate $n^{-1 / 5}$. Also, for certain positive constants $C_{1}, C_{2}, \sigma_{1}, \sigma_{2}$,
\[

$$
\begin{aligned}
n^{3 / 10}\left(\widehat{h}-\widehat{h}_{0}\right) \rightsquigarrow N\left(0, \sigma_{1}^{2}\right), & n\left(L(\widehat{h})-L\left(\widehat{h}_{0}\right) \rightsquigarrow C_{1} \chi_{1}^{2}\right. \\
n^{3 / 10}\left(h_{0}-\widehat{h}_{0}\right) \rightsquigarrow N\left(0, \sigma_{2}^{2}\right), & n\left(L\left(h_{0}\right)-L\left(\widehat{h}_{0}\right) \rightsquigarrow C_{2} \chi_{1}^{2} .\right.
\end{aligned}
$$
\]

Thus, the relative rate of convergence of $\widehat{h}$ is

$$
\frac{\widehat{h}-\widehat{h}_{0}}{\widehat{h}_{0}}=O_{P}\left(\frac{n^{3 / 10}}{n^{1 / 5}}\right)=O_{P}\left(n^{-1 / 10}\right) .
$$

This slow rate shows that estimating the bandwidth is difficult. This rate is intrinsic to the problem of bandwidth selection since it is also true that

$$
\frac{\widehat{h}_{0}-h_{0}}{h_{0}}=O_{P}\left(\frac{n^{3 / 10}}{n^{1 / 5}}\right)=O_{P}\left(n^{-1 / 10}\right)
$$

### 5.4 Local Regression

Now we turn to local nonparametric regression. Suppose that $x_{i} \in \mathbb{R}$ is scalar and consider the regression model (5.1). In this section we consider estimators of $r(x)$ obtained by taking a weighted average of the $Y_{i} \mathrm{~s}$, giving higher weight to those points near $x$. We begin with the kernel regression estimator.
5.39 Definition. Let $h>0$ be a positive number, called the bandwidth. The Nadaraya-Watson kernel estimator is defined by

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i} \tag{5.40}
\end{equation*}
$$

where $K$ is a kernel (as defined in Section 4.2) and the weights $\ell_{i}(x)$ are given by

$$
\begin{equation*}
\ell_{i}(x)=\frac{K\left(\frac{x-x_{i}}{h}\right)}{\sum_{j=1}^{n} K\left(\frac{x-x_{j}}{h}\right)} \tag{5.41}
\end{equation*}
$$

5.42 Remark. The local average estimator in Example 5.26 is a kernel estimator based on the boxcar kernel.
5.43 Example (CMB data). Recall the CMB data from Figure 5.1. Figure 5.2 shows four different kernel regression fits (using just the first 400 data points) based on increasing bandwidths. The top two plots are based on small
bandwidths and the fits are too rough. The bottom right plot is based on large bandwidth and the fit is too smooth. The bottom left plot is just right. The bottom right plot also shows the presence of bias near the boundaries. As we shall see, this is a general feature of kernel regression. The bottom plot in Figure 5.3 shows a kernel fit to all the data points. The bandwidth was chosen by cross-validation.


FIGURE 5.2. Four kernel regressions for the CMB data using just the first 400 data points. The bandwidths used were $h=1$ (top left), $h=10$ (top right), $h=50$ (bottom left), $h=200$ (bottom right). As the bandwidth $h$ increases, the estimated function goes from being too rough to too smooth.

The choice of kernel $K$ is not too important. Estimates obtained by using different kernels are usually numerically very similar. This observation is confirmed by theoretical calculations which show that the risk is very insensitive to the choice of kernel; see Section 6.2.3 of Scott (1992). We will often use the tricube kernel in examples. What does matter much more is the choice of bandwidth $h$ which controls the amount of smoothing. Small bandwidths
give very rough estimates while larger bandwidths give smoother estimates. In general, we will let the bandwidth depend on the sample size so we sometimes write $h_{n}$.

The following theorem shows how the bandwidth affects the estimator. To state these results we need to make some assumption about the behavior of $x_{1}, \ldots, x_{n}$ as $n$ increases. For the purposes of the theorem, we will assume that these are random draws from some density $f$.
5.44 Theorem. The risk (using integrated squared error loss) of the NadarayaWatson kernel estimator is

$$
\begin{align*}
R\left(\widehat{r}_{n}, r\right)= & \frac{h_{n}^{4}}{4}\left(\int x^{2} K(x) d x\right)^{2} \int\left(r^{\prime \prime}(x)+2 r^{\prime}(x) \frac{f^{\prime}(x)}{f(x)}\right)^{2} d x \\
& +\frac{\sigma^{2} \int K^{2}(x) d x}{n h_{n}} \int \frac{1}{f(x)} d x+o\left(n h_{n}^{-1}\right)+o\left(h_{n}^{4}\right) \tag{5.45}
\end{align*}
$$

as $h_{n} \rightarrow 0$ and $n h_{n} \rightarrow \infty$.
The first term in (5.45) is the squared bias and the second term is the variance. What is especially notable is the presence of the term

$$
\begin{equation*}
2 r^{\prime}(x) \frac{f^{\prime}(x)}{f(x)} \tag{5.46}
\end{equation*}
$$

in the bias. We call (5.46) the design bias since it depends on the design, that is, the distribution of the $x_{i}$ 's. This means that the bias is sensitive to the position of the $x_{i}$ s. Furthermore, it can be shown that kernel estimators also have high bias near the boundaries. This is known as boundary bias. We will see that we can reduce these biases by using a refinement called local polynomial regression.

If we differentiate (5.45) and set the result equal to 0 , we find that the optimal bandwidth $h_{*}$ is

$$
\begin{equation*}
h_{*}=\left(\frac{1}{n}\right)^{1 / 5}\left(\frac{\sigma^{2} \int K^{2}(x) d x \int d x / f(x)}{\left(\int x^{2} K^{2}(x) d x\right)^{2} \int\left(r^{\prime \prime}(x)+2 r^{\prime}(x) \frac{f^{\prime}(x)}{f(x)}\right)^{2} d x}\right)^{1 / 5} . \tag{5.47}
\end{equation*}
$$

Thus, $h_{*}=O\left(n^{-1 / 5}\right)$. Plugging $h_{*}$ back into (5.45) we see that the risk decreases at rate $O\left(n^{-4 / 5}\right)$. In (most) parametric models, the risk of the maximum likelihood estimator decreases to 0 at rate $1 / n$. The slower rate $n^{-4 / 5}$ is the price of using nonparametric methods. In practice, we cannot use the bandwidth given in (5.47) since $h_{*}$ depends on the unknown function $r$. Instead, we use leave-one-out cross-validation as described in Theorem 5.34.
5.48 Example. Figure 5.3 shows the cross-validation score for the CMB example as a function of the effective degrees of freedom. The optimal smoothing parameter was chosen to minimize this score. The resulting fit is also shown in the figure. Note that the fit gets quite variable to the right. Later we will deal with the nonconstant variance and we will add confidence bands to the fit.


FIGURE 5.3. Top: The cross-validation (CV) score as a function of the effective degrees of freedom. Bottom: the kernel fit using the bandwidth that minimizes the cross-validation score.

Local Polynomials. Kernel estimators suffer from boundary bias and design bias. These problems can be alleviated by using a generalization of kernel regression called local polynomial regression.

To motivate this estimator, first consider choosing an estimator $a \equiv \widehat{r}_{n}(x)$ to minimize the sums of squares $\sum_{i=1}^{n}\left(Y_{i}-a\right)^{2}$. The solution is the constant
function $\widehat{r}_{n}(x)=\bar{Y}$ which is obviously not a good estimator of $r(x)$. Now define the weight function $w_{i}(x)=K\left(\left(x_{i}-x\right) / h\right)$ and choose $a \equiv \widehat{r}_{n}(x)$ to minimize the weighted sums of squares

$$
\begin{equation*}
\sum_{i=1}^{n} w_{i}(x)\left(Y_{i}-a\right)^{2} \tag{5.49}
\end{equation*}
$$

From elementary calculus, we see that the solution is

$$
\widehat{r}_{n}(x) \equiv \frac{\sum_{i=1}^{n} w_{i}(x) Y_{i}}{\sum_{i=1}^{n} w_{i}(x)}
$$

which is exactly the kernel regression estimator. This gives us an interesting interpretation of the kernel estimator: it is a locally constant estimator, obtained from locally weighted least squares.

This suggests that we might improve the estimator by using a local polynomial of degree $p$ instead of a local constant. Let $x$ be some fixed value at which we want to estimate $r(x)$. For values $u$ in a neighborhood of $x$, define the polynomial

$$
\begin{equation*}
P_{x}(u ; a)=a_{0}+a_{1}(u-x)+\frac{a_{2}}{2!}(u-x)^{2}+\cdots+\frac{a_{p}}{p!}(u-x)^{p} . \tag{5.50}
\end{equation*}
$$

We can approximate a smooth regression function $r(u)$ in a neighborhood of the target value $x$ by the polynomial:

$$
\begin{equation*}
r(u) \approx P_{x}(u ; a) \tag{5.51}
\end{equation*}
$$

We estimate $a=\left(a_{0}, \ldots, a_{p}\right)^{T}$ by choosing $\widehat{a}=\left(\widehat{a}_{0}, \ldots, \widehat{a}_{p}\right)^{T}$ to minimize the locally weighted sums of squares

$$
\begin{equation*}
\sum_{i=1}^{n} w_{i}(x)\left(Y_{i}-P_{x}\left(X_{i} ; a\right)\right)^{2} \tag{5.52}
\end{equation*}
$$

The estimator $\widehat{a}$ depends on the target value $x$ so we write $\widehat{a}(x)=\left(\widehat{a}_{0}(x), \ldots\right.$ , $\left.\widehat{a}_{p}(x)\right)^{T}$ if we want to make this dependence explicit. The local estimate of $r$ is

$$
\widehat{r}_{n}(u)=P_{x}(u ; \widehat{a}) .
$$

In particular, at the target value $u=x$ we have

$$
\begin{equation*}
\widehat{r}_{n}(x)=P_{x}(x ; \widehat{a})=\widehat{a}_{0}(x) . \tag{5.53}
\end{equation*}
$$

Warning! Although $\widehat{r}_{n}(x)$ only depends on $\widehat{a}_{0}(x)$, this is not equivalent to simply fitting a local constant.

Setting $p=0$ gives back the kernel estimator. The special case where $p=1$ is called local linear regression and this is the version we recommend as a default choice. As we shall see, local polynomial estimators, and in particular local linear estimators, have some remarkable properties as shown by Fan (1992) and Hastie and Loader (1993). Many of the results that follow are from those papers.

To find $\widehat{a}(x)$, it is helpful to re-express the problem in vector notation. Let

$$
X_{x}=\left(\begin{array}{llll}
1 & x_{1}-x & \cdots & \frac{\left(x_{1}-x\right)^{p}}{p!}  \tag{5.54}\\
1 & x_{2}-x & \cdots & \frac{\left(x_{2}-x\right)^{p}}{p!} \\
\vdots & \vdots & \ddots & \vdots \\
1 & x_{n}-x & \cdots & \frac{\left(x_{n}-x\right)^{p}}{p!}
\end{array}\right)
$$

and let $W_{x}$ be the $n \times n$ diagonal matrix whose ( $i, i$ ) component is $w_{i}(x)$. We can rewrite (5.52) as

$$
\begin{equation*}
\left(Y-X_{x} a\right)^{T} W_{x}\left(Y-X_{x} a\right) . \tag{5.55}
\end{equation*}
$$

Minimizing (5.55) gives the weighted least squares estimator

$$
\begin{equation*}
\widehat{a}(x)=\left(X_{x}^{T} W_{x} X_{x}\right)^{-1} X_{x}^{T} W_{x} Y \tag{5.56}
\end{equation*}
$$

In particular, $\widehat{r}_{n}(x)=\widehat{a}_{0}(x)$ is the inner product of the first row of $\left(X_{x}^{T} W_{x} X_{x}\right)^{-1} X_{x}^{T} W_{x}$ with $Y$. Thus we have:
5.57 Theorem. The local polynomial regression estimate is

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i} \tag{5.58}
\end{equation*}
$$

where $\ell(x)^{T}=\left(\ell_{1}(x), \ldots, \ell_{n}(x)\right)$,

$$
\ell(x)^{T}=e_{1}^{T}\left(X_{x}^{T} W_{x} X_{x}\right)^{-1} X_{x}^{T} W_{x}
$$

$e_{1}=(1,0, \ldots, 0)^{T}$ and $X_{x}$ and $W_{x}$ are defined in (5.54). This estimator has mean

$$
\mathbb{E}\left(\widehat{r}_{n}(x)\right)=\sum_{i=1}^{n} \ell_{i}(x) r\left(x_{i}\right)
$$

and variance

$$
\mathbb{V}\left(\widehat{r}_{n}(x)\right)=\sigma^{2} \sum_{i=1}^{n} \ell_{i}(x)^{2}=\sigma^{2}\|\ell(x)\|^{2}
$$

Once again, our estimate is a linear smoother and we can choose the bandwidth by minimizing the cross-validation formula given in Theorem 5.34.
5.59 Example (LIDAR). These data were introduced in Example 4.5. Figure 5.4 shows the 221 observations. The top left plot shows the data and the fitted function using local linear regression. The cross-validation curve (not shown) has a well-defined minimum at $h \approx 37$ corresponding to 9 effective degrees of freedom. The fitted function uses this bandwidth. The top right plot shows the residuals. There is clear heteroscedasticity (nonconstant variance). The bottom left plot shows the estimate of $\sigma(x)$ using the method in Section 5.6. (with $h=146$ chosen via cross-validation). Next we use the method in Section 5.7 to compute 95 percent confidence bands. The resulting bands are shown in the lower right plot. As expected, there is much greater uncertainty for larger values of the covariate.

## Local Linear Smoothing

5.60 Theorem. When $p=1, \widehat{r}_{n}(x)=\sum_{i=1}^{n} \ell_{i}(x) Y_{i}$ where

$$
\begin{gather*}
\ell_{i}(x)=\frac{b_{i}(x)}{\sum_{j=1}^{n} b_{j}(x)}, \\
b_{i}(x)=K\left(\frac{x_{i}-x}{h}\right)\left(S_{n, 2}(x)-\left(x_{i}-x\right) S_{n, 1}(x)\right) \tag{5.61}
\end{gather*}
$$

and

$$
S_{n, j}(x)=\sum_{i=1}^{n} K\left(\frac{x_{i}-x}{h}\right)\left(x_{i}-x\right)^{j}, \quad j=1,2 .
$$

5.62 Example. Figure 5.5 shows the local regression for the CMB data for $p=0$ and $p=1$. The bottom plots zoom in on the left boundary. Note that for $p=0$ (the kernel estimator), the fit is poor near the boundaries due to boundary bias.
5.63 Example (Doppler function). Let

$$
\begin{equation*}
r(x)=\sqrt{x(1-x)} \sin \left(\frac{2.1 \pi}{x+.05}\right), \quad 0 \leq x \leq 1 \tag{5.64}
\end{equation*}
$$

which is called the Doppler function. This function is difficult to estimate and provides a good test case for nonparametric regression methods. The function is spatially inhomogeneous which means that its smoothness (second


FIGURE 5.4. The LIDAR data from Example 5.59. Top left: data and the fitted function using local linear regression with $h \approx 37$ (chosen by cross-validation). Top right: the residuals. Bottom left: estimate of $\sigma(x)$. Bottom right: 95 percent confidence bands.
derivative) varies over $x$. The function is plotted in the top left plot of Figure 5.6. The top right plot shows 1000 data points simulated from $Y_{i}=r(i / n)+\sigma \epsilon_{i}$ with $\sigma=.1$ and $\epsilon_{i} \sim N(0,1)$. The bottom left plot shows the cross-validation score versus the effective degrees of freedom using local linear regression. The minimum occurred at 166 degrees of freedom corresponding to a bandwidth of .005 . The fitted function is shown in the bottom right plot. The fit has high effective degrees of freedom and hence the fitted function is very wiggly. This is because the estimate is trying to fit the rapid fluctuations of the function near $x=0$. If we used more smoothing, the right-hand side of the fit would look better at the cost of missing the structure near $x=0$. This is always a problem when estimating spatially inhomogeneous functions. We'll discuss that further when we talk about wavelets in Chapter 9.


FIGURE 5.5. Locally weighted regressions using local polynomials of order $p=0$ (top left) and $p=1$ (top right). The bottom plots show the left boundary in more detail ( $p=0$ bottom left and $p=1$ bottom right). Notice that the boundary bias is reduced by using local linear estimation ( $p=1$ ).

The following theorem gives the large sample behavior of the risk of the local linear estimator and shows why local linear regression is better than kernel regression. A proof can be found in Fan (1992) and Fan and Gijbels (1996).
5.65 Theorem. Let $Y_{i}=r\left(X_{i}\right)+\sigma\left(X_{i}\right) \epsilon_{i}$ for $i=1, \ldots, n$ and $a \leq X_{i} \leq b$. Assume that $X_{1}, \ldots, X_{n}$ are a sample from a distribution with density $f$ and that (i) $f(x)>0$, (ii) $f, r^{\prime \prime}$ and $\sigma^{2}$ are continuous in a neighborhood of $x$, and (iii) $h_{n} \rightarrow 0$ and $n h_{n} \rightarrow \infty$. Let $x \in(a, b)$. Given $X_{1}, \ldots, X_{n}$, we have the following: the local linear estimator and the kernel estimator both have variance

$$
\begin{equation*}
\frac{\sigma^{2}(x)}{f(x) n h_{n}} \int K^{2}(u) d u+o_{P}\left(\frac{1}{n h_{n}}\right) . \tag{5.66}
\end{equation*}
$$



FIGURE 5.6. The Doppler function estimated by local linear regression. The function (top left), the data (top right), the cross-validation score versus effective degrees of freedom (bottom left), and the fitted function (bottom right).

The Nadaraya-Watson kernel estimator has bias

$$
\begin{equation*}
h_{n}^{2}\left(\frac{1}{2} r^{\prime \prime}(x)+\frac{r^{\prime}(x) f^{\prime}(x)}{f(x)}\right) \int u^{2} K(u) d u+o_{P}\left(h^{2}\right) \tag{5.67}
\end{equation*}
$$

whereas the local linear estimator has asymptotic bias

$$
\begin{equation*}
h_{n}^{2} \frac{1}{2} r^{\prime \prime}(x) \int u^{2} K(u) d u+o_{P}\left(h^{2}\right) \tag{5.68}
\end{equation*}
$$

Thus, the local linear estimator is free from design bias. At the boundary points $a$ and b, the Nadaraya-Watson kernel estimator has asymptotic bias of order $h_{n}$ while the local linear estimator has bias of order $h_{n}^{2}$. In this sense, local linear estimation eliminates boundary bias.
5.69 Remark. The above result holds more generally for local polynomials of order $p$. Generally, taking $p$ odd reduces design bias and boundary bias without increasing variance.

### 5.5 Penalized Regression, Regularization and Splines

Consider once again the regression model

$$
Y_{i}=r\left(x_{i}\right)+\epsilon_{i}
$$

and suppose we estimate $r$ by choosing $\widehat{r}_{n}(x)$ to minimize the sums of squares

$$
\sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}
$$

Minimizing over all linear functions (i.e., functions of the form $\beta_{0}+\beta_{1} x$ ) yields the least squares estimator. Minimizing over all functions yields a function that interpolates the data. In the previous section we avoided these two extreme solutions by replacing the sums of squares with a locally weighted sums of squares. An alternative way to get solutions in between these extremes is to minimize the penalized sums of squares

$$
\begin{equation*}
M(\lambda)=\sum_{i}\left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}+\lambda J(r) \tag{5.70}
\end{equation*}
$$

where $J(r)$ is some roughness penalty. Adding a penalty term to the criterion we are optimizing is sometimes called regularization.

We will focus on the special case

$$
\begin{equation*}
J(r)=\int\left(r^{\prime \prime}(x)\right)^{2} d x \tag{5.71}
\end{equation*}
$$

The parameter $\lambda$ controls the trade-off between fit (the first term of 5.70) and the penalty. Let $\widehat{r}_{n}$ denote the function that minimizes $M(\lambda)$. When $\lambda=0$, the solution is the interpolating function. When $\lambda \rightarrow \infty, \widehat{r}_{n}$ converges to the least squares line. The parameter $\lambda$ controls the amount of smoothing. What does $\widehat{r}_{n}$ look like for $0<\lambda<\infty$ ? To answer this question, we need to define splines.

A spline is a special piecewise polynomial. ${ }^{4}$ The most commonly used splines are piecewise cubic splines.
5.72 Definition. Let $\xi_{1}<\xi_{2}<\cdots<\xi_{k}$ be a set of ordered points-called knots - contained in some interval $(a, b)$. A cubic spline is a continuous function $r$ such that (i) $r$ is a cubic polynomial over $\left(\xi_{1}, \xi_{2}\right)$, ... and (ii) $r$ has continuous first and second derivatives at the knots. More generally,

[^5]an $M^{\text {th }}$-order spline is a piecewise $M-1$ degree polynomial with $M-2$ continuous derivatives at the knots. A spline that is linear beyond the boundary knots is called a natural spline.

Cubic splines $(M=4)$ are the most common splines used in practice. They arise naturally in the penalized regression framework as the following theorem shows.
5.73 Theorem. The function $\widehat{r}_{n}(x)$ that minimizes $M(\lambda)$ with penalty (5.71) is a natural cubic spline with knots at the data points. The estimator $\widehat{r}_{n}$ is called $a$ smoothing spline.

The theorem above does not give an explicit form for $\widehat{r}_{n}$. To do so, we will construct a basis for the set of splines.
5.74 Theorem. Let $\xi_{1}<\xi_{2}<\cdots<\xi_{k}$ be knots contained in an interval $(a, b)$. Define $h_{1}(x)=1, h_{2}(x)=x, h_{3}(x)=x^{2}, h_{4}(x)=x^{3}, h_{j}(x)=\left(x-\xi_{j-4}\right)_{+}^{3}$ for $j=5, \ldots, k+4$. The functions $\left\{h_{1}, \ldots, h_{k+4}\right\}$ form a basis for the set of cubic splines at these knots, called the truncated power basis. Thus, any cubic spline $r(x)$ with these knots can be written as

$$
\begin{equation*}
r(x)=\sum_{j=1}^{k+4} \beta_{j} h_{j}(x) \tag{5.75}
\end{equation*}
$$

We now introduce a different basis for the set of natural splines called the B-spline basis that is particularly well suited for computation. These are defined as follows.

Let $\xi_{0}=a$ and $\xi_{k+1}=b$. Define new knots $\tau_{1}, \ldots, \tau_{M}$ such that

$$
\tau_{1} \leq \tau_{2} \leq \tau_{3} \leq \cdots \leq \tau_{M} \leq \xi_{0}
$$

$\tau_{j+M}=\xi_{j}$ for $j=1, \ldots, k$, and

$$
\xi_{k+1} \leq \tau_{k+M+1} \leq \cdots \leq \tau_{k+2 M}
$$

The choice of extra knots is arbitrary; usually one takes $\tau_{1}=\cdots=\tau_{M}=\xi_{0}$ and $\xi_{k+1}=\tau_{k+M+1}=\cdots=\tau_{k+2 M}$. We define the basis functions recursively as follows. First we define

$$
B_{i, 1}= \begin{cases}1 & \text { if } \tau_{i} \leq x<\tau_{i+1} \\ 0 & \text { otherwise }\end{cases}
$$

for $i=1, \ldots, k+2 M-1$. Next, for $m \leq M$ we define

$$
B_{i, m}=\frac{x-\tau_{i}}{\tau_{i+m-1}-\tau_{i}} B_{i, m-1}+\frac{\tau_{i+m}-x}{\tau_{i+m}-\tau_{i+1}} B_{i+1, m-1}
$$



FIGURE 5.7. Cubic B-spline basis using nine equally spaced knots on $(0,1)$.
for $i=1, \ldots, k+2 M-m$. It is understood that if the denominator is 0 , then the function is defined to be 0 .
5.76 Theorem. The functions $\left\{B_{i, 4}, i=1, \ldots, k+4\right\}$ are a basis for the set of cubic splines. They are called the $\mathbf{B}$-spline basis functions.

The advantage of the B-spline basis functions is that they have compact support which makes it possible to speed up calculations. See Hastie et al. (2001) for details. Figure 5.7 shows the cubic B-spline basis using nine equally spaced knots on $(0,1)$.

We are now in a position to describe the spline estimator in more detail. According to Theorem 5.73, $\widehat{r}$ is a natural cubic spline. Hence, we can write

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{j=1}^{N} \widehat{\beta}_{j} B_{j}(x) \tag{5.77}
\end{equation*}
$$

where $B_{1}, \ldots, B_{N}$ are a basis for the natural splines (such as the B-splines with $N=n+4)$. Thus, we only need to find the coefficients $\widehat{\beta}=\left(\widehat{\beta}_{1}, \ldots, \widehat{\beta}_{N}\right)^{T}$. By expanding $r$ in the basis we can now rewrite the minimization as follows:

$$
\begin{equation*}
\text { minimize : } \quad(Y-B \beta)^{T}(Y-B \beta)+\lambda \beta^{T} \Omega \beta \tag{5.78}
\end{equation*}
$$

where $B_{i j}=B_{j}\left(X_{i}\right)$ and $\Omega_{j k}=\int B_{j}^{\prime \prime}(x) B_{k}^{\prime \prime}(x) d x$.
5.79 Theorem. The value of $\beta$ that minimizes (5.78) is ${ }^{5}$

$$
\begin{equation*}
\widehat{\beta}=\left(B^{T} B+\lambda \Omega\right)^{-1} B^{T} Y . \tag{5.80}
\end{equation*}
$$

Splines are another example of linear smoothers.
5.81 Theorem. The smoothing spline $\widehat{r}_{n}(x)$ is a linear smoother, that is, there exist weights $\ell(x)$ such that $\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x)$. In particular, the smoothing matrix $L$ is

$$
\begin{equation*}
L=B\left(B^{T} B+\lambda \Omega\right)^{-1} B^{T} \tag{5.82}
\end{equation*}
$$

and the vector $\mathbf{r}$ of fitted values is given by

$$
\begin{equation*}
\mathbf{r}=L Y \tag{5.83}
\end{equation*}
$$

If we had done ordinary linear regression of $Y$ on $B$, the hat matrix would be $L=B\left(B^{T} B\right)^{-1} B^{T}$ and the fitted values would interpolate the observed data. The effect of the term $\lambda \Omega$ in (5.82) is to shrink the regression coefficients towards a subspace, which results in a smoother fit. As before, we define the effective degrees of freedom by $\nu=\operatorname{tr}(L)$ and we choose the smoothing parameter $\lambda$ by minimizing either the cross-validation score (5.35) or the generalized cross-validation score (5.36).
5.84 Example. Figure 5.8 shows the smoothing spline with cross-validation for the CMB data. The effective number of degrees of freedom is 8.8 . The fit is smoother than the local regression estimator. This is certainly visually more appealing, but the difference between the two fits is small compared to the width of the confidence bands that we will compute later.

Silverman (1984) proved that spline estimates $\widehat{r}_{n}(x)$ are approximately kernel estimates in the sense that

$$
\ell_{i}(x) \approx \frac{1}{f\left(x_{i}\right) h\left(x_{i}\right)} K\left(\frac{x_{i}-x}{h\left(x_{i}\right)}\right)
$$

where $f(x)$ is the density of the covariate (treated here as random),

$$
h(x)=\left(\frac{\lambda}{n f(x)}\right)^{1 / 4}
$$

and

$$
K(t)=\frac{1}{2} \exp \left\{-\frac{|t|}{\sqrt{2}}\right\} \sin \left(\frac{|t|}{\sqrt{2}}+\frac{\pi}{4}\right) .
$$

[^6]

FIGURE 5.8. Smoothing spline for the CMB data. The smoothing parameter was chosen by cross-validation.

Another nonparametric method that uses splines is called the regression spline method. Rather than placing a knot at each data point, we instead use fewer knots. We then do ordinary linear regression on the basis matrix $B$ with no regularization. The fitted values for this estimator are $\mathbf{r}=L Y$ with $L=B\left(B^{T} B\right)^{-1} B^{T}$. The difference between this estimate and (5.82) is that the basis matrix $B$ is based on fewer knots and there is no shrinkage factor $\lambda \Omega$. The amount of smoothing is instead controlled by the choice of the number (and placement) of the knots. By using fewer knots, one can save computation time. These spline methods are discussed in detail in Ruppert et al. (2003).

### 5.6 Variance Estimation

Next we consider several methods for estimating $\sigma^{2}$. For linear smoothers, there is a simple, nearly unbiased estimate of $\sigma^{2}$.
5.85 Theorem. Let $\widehat{r}_{n}(x)$ be a linear smoother. Let

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{\sum_{i=1}^{n}\left(Y_{i}-\widehat{r}\left(x_{i}\right)\right)^{2}}{n-2 \nu+\widetilde{\nu}} \tag{5.86}
\end{equation*}
$$

where

$$
\nu=\operatorname{tr}(L), \quad \widetilde{\nu}=\operatorname{tr}\left(L^{T} L\right)=\sum_{i=1}^{n}\left\|\ell\left(x_{i}\right)\right\|^{2}
$$

If $r$ is sufficiently smooth, $\nu=o(n)$ and $\widetilde{\nu}=o(n)$ then $\widehat{\sigma}^{2}$ is a consistent estimator of $\sigma^{2}$.

We will now outline the proof of this result. Recall that if $Y$ is a random vector and $Q$ is a symmetric matrix then $Y^{T} Q Y$ is called a quadratic form and it is well known that

$$
\begin{equation*}
\mathbb{E}\left(Y^{T} Q Y\right)=\operatorname{tr}(Q V)+\mu^{T} Q \mu \tag{5.87}
\end{equation*}
$$

where $V=\mathbb{V}(Y)$ is the covariance matrix of $Y$ and $\mu=\mathbb{E}(Y)$ is the mean vector. Now,

$$
Y-\mathbf{r}=Y-L Y=(I-L) Y
$$

and so

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{Y^{T} \Lambda Y}{\operatorname{tr}(\Lambda)} \tag{5.88}
\end{equation*}
$$

where $\Lambda=(I-L)^{T}(I-L)$. Hence,

$$
\mathbb{E}\left(\widehat{\sigma}^{2}\right)=\frac{\mathbb{E}\left(Y^{T} \Lambda Y\right)}{\operatorname{tr}(\Lambda)}=\sigma^{2}+\frac{\mathbf{r}^{T} \Lambda \mathbf{r}}{n-2 \nu+\widetilde{\nu}}
$$

Assuming that $\nu$ and $\widetilde{\nu}$ do not grow too quickly, and that $r$ is smooth, the last term is small for large $n$ and hence $\mathbb{E}\left(\widehat{\sigma}^{2}\right) \approx \sigma^{2}$. Similarly, one can show that $\mathbb{V}\left(\widehat{\sigma}^{2}\right) \rightarrow 0$.

Here is another estimator, due to Rice (1984). Suppose that the $x_{i} \mathrm{~S}$ are ordered. Define

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{1}{2(n-1)} \sum_{i=1}^{n-1}\left(Y_{i+1}-Y_{i}\right)^{2} \tag{5.89}
\end{equation*}
$$

The motivation for this estimator is as follows. Assuming $r(x)$ is smooth, we have $r\left(x_{i+1}\right)-r\left(x_{i}\right) \approx 0$ and hence

$$
Y_{i+1}-Y_{i}=\left[r\left(x_{i+1}\right)+\epsilon_{i+1}\right]-\left[r\left(x_{i}\right)+\epsilon_{i}\right] \approx \epsilon_{i+1}-\epsilon_{i}
$$

and hence $\left(Y_{i+1}-Y_{i}\right)^{2} \approx \epsilon_{i+1}^{2}+\epsilon_{i}^{2}-2 \epsilon_{i+1} \epsilon_{i}$. Therefore,

$$
\begin{align*}
\mathbb{E}\left(Y_{i+1}-Y_{i}\right)^{2} & \approx \mathbb{E}\left(\epsilon_{i+1}^{2}\right)+\mathbb{E}\left(\epsilon_{i}^{2}\right)-2 \mathbb{E}\left(\epsilon_{i+1}\right) \mathbb{E}\left(\epsilon_{i}\right) \\
& =\mathbb{E}\left(\epsilon_{i+1}^{2}\right)+\mathbb{E}\left(\epsilon_{i}^{2}\right)=2 \sigma^{2} . \tag{5.90}
\end{align*}
$$

Thus, $\mathbb{E}\left(\widehat{\sigma}^{2}\right) \approx \sigma^{2}$. A variation of this estimator, due to Gasser et al. (1986) is

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{1}{n-2} \sum_{i=2}^{n-1} c_{i}^{2} \delta_{i}^{2} \tag{5.91}
\end{equation*}
$$

where

$$
\begin{array}{ll}
\delta_{i}=a_{i} Y_{i-1}+b_{i} Y_{i+1}-Y_{i}, & a_{i}=\left(x_{i+1}-x_{i}\right) /\left(x_{i+1}-x_{i-1}\right), \\
b_{i}=\left(x_{i}-x_{i-1}\right) /\left(x_{i+1}-x_{i-1}\right), & c_{i}^{2}=\left(a_{i}^{2}+b_{i}^{2}+1\right)^{-1} .
\end{array}
$$

The intuition of this estimator is that it is the average of the residuals that result from fitting a line to the first and third point of each consecutive triple of design points.
5.92 Example. The variance looks roughly constant for the first 400 observations of the CMB data. Using a local linear fit, we applied the two variance estimators. Equation (5.86) yields $\widehat{\sigma}^{2}=408.29$ while equation (5.89) yields $\widehat{\sigma}^{2}=394.55$.

So far we have assumed homoscedasticity meaning that $\sigma^{2}=\mathbb{V}\left(\epsilon_{i}\right)$ does not vary with $x$. In the CMB example this is blatantly false. Clearly, $\sigma^{2}$ increases with $x$ so the data are heteroscedastic. The function estimate $\widehat{r}_{n}(x)$ is relatively insensitive to heteroscedasticity. However, when it comes to making confidence bands for $r(x)$, we must take into account the nonconstant variance.

We will take the following approach. See Yu and Jones (2004) and references therein for other approaches. Suppose that

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\sigma\left(x_{i}\right) \epsilon_{i} \tag{5.93}
\end{equation*}
$$

Let $Z_{i}=\log \left(Y_{i}-r\left(x_{i}\right)\right)^{2}$ and $\delta_{i}=\log \epsilon_{i}^{2}$. Then,

$$
\begin{equation*}
Z_{i}=\log \left(\sigma^{2}\left(x_{i}\right)\right)+\delta_{i} \tag{5.94}
\end{equation*}
$$

This suggests estimating $\log \sigma^{2}(x)$ by regressing the $\log$ squared residuals on $x$. We proceed as follows.

## Variance Function Estimation

1. Estimate $r(x)$ with any nonparametric method to get an estimate

$$
\widehat{r}_{n}(x)
$$

2. Define $Z_{i}=\log \left(Y_{i}-\widehat{r}_{n}\left(x_{i}\right)\right)^{2}$.
3. Regress the $Z_{i}$ 's on the $x_{i}$ 's (again using any nonparametric method) to get an estimate $\widehat{q}(x)$ of $\log \sigma^{2}(x)$ and let

$$
\begin{equation*}
\widehat{\sigma}^{2}(x)=e^{\widehat{q}(x)} \tag{5.95}
\end{equation*}
$$

5.96 Example. The solid line in Figure 5.9 shows the $\log$ of $\widehat{\sigma}^{2}(x)$ for the CMB example. I used local linear estimation and I used cross-validation to choose the bandwidth. The estimated optimal bandwidth for $\widehat{r}_{n}$ was $h=42$ while the estimated optimal bandwidth for the $\log$ variance was $h=160$. In this example, there turns out to be an independent estimate of $\sigma(x)$. Specifically, because the physics of the measurement process is well understood, physicists can compute a reasonably accurate approximation to $\sigma^{2}(x)$. The $\log$ of this function is the dotted line on the plot.


FIGURE 5.9. The dots are the log squared residuals. The solid line shows the log of the estimated standard variance $\widehat{\sigma}^{2}(x)$ as a function of $x$. The dotted line shows the $\log$ of the true $\sigma^{2}(x)$ which is known (to reasonable accuracy) through prior knowledge.

A drawback of this approach is that the log of a very small residual will be a large outlier. An alternative is to directly smooth the squared residuals. In
this case, one might fit a model of the type discussed in Section 5.10 since the model does not have an additive form and the error will not be Normal.

### 5.7 Confidence Bands

In this section we will construct confidence bands for $r(x)$. Typically these bands are of the form

$$
\begin{equation*}
\widehat{r}_{n}(x) \pm c \operatorname{se}(x) \tag{5.97}
\end{equation*}
$$

where $\operatorname{se}(x)$ is an estimate of the standard deviation of $\widehat{r}_{n}(x)$ and $c>0$ is some constant. Before we proceed, we discuss a pernicious problem that arises whenever we do smoothing, namely, the bias problem.

The Bias Problem. Confidence bands like those in (5.97), are not really confidence bands for $r(x)$, rather, they are confidence bands for $\bar{r}_{n}(x)=\mathbb{E}\left(\widehat{r}_{n}(x)\right)$ which you can think of as a smoothed version of $r(x)$. Getting a confidence set for the true function $r(x)$ is complicated for reasons we now explain.

Denote the mean and standard deviation of $\widehat{r}_{n}(x)$ by $\bar{r}_{n}(x)$ and $s_{n}(x)$. Then,

$$
\begin{aligned}
\frac{\widehat{r}_{n}(x)-r(x)}{s_{n}(x)} & =\frac{\widehat{r}_{n}(x)-\bar{r}_{n}(x)}{s_{n}(x)}+\frac{\bar{r}_{n}(x)-r(x)}{s_{n}(x)} \\
& =Z_{n}(x)+\frac{\operatorname{bias}\left(\widehat{r}_{n}(x)\right)}{\sqrt{\text { variance }\left(\widehat{r}_{n}(x)\right)}}
\end{aligned}
$$

where $Z_{n}(x)=\left(\widehat{r}_{n}(x)-\bar{r}_{n}(x)\right) / s_{n}(x)$. Typically, the first term $Z_{n}(x)$ converges to a standard Normal from which one derives confidence bands. The second term is the bias divided by the standard deviation. In parametric inference, the bias is usually smaller than the standard deviation of the estimator so this term goes to zero as the sample size increases. In nonparametric inference, we have seen that optimal smoothing corresponds to balancing the bias and the standard deviation. The second term does not vanish even with large sample sizes.

> The presence of this second, nonvanishing term introduces a bias into the Normal limit. The result is that the confidence interval will not be centered around the true function $r$ due to the smoothing bias $\bar{r}_{n}(x)-r(x)$.

There are several things we can do about this problem. The first is: live with it. In other words, just accept the fact that the confidence band is for $\bar{r}_{n}$ not $r$.

There is nothing wrong with this as long as we are careful when we report the results to make it clear that the inferences are for $\bar{r}_{n}$ not $r$. A second approach is to estimate the bias function $\bar{r}_{n}(x)-r(x)$. This is difficult to do. Indeed, the leading term of the bias is $r^{\prime \prime}(x)$ and estimating the second derivative of $r$ is much harder than estimating $r$. This requires introducing extra smoothness conditions which then bring into question the original estimator that did not use this extra smoothness. This has a certain unpleasant circularity to it. ${ }^{6}$ A third approach is to undersmooth. If we smooth less than the optimal amount then the bias will decrease asymptotically relative to the variance. Unfortunately, there does not seem to be a simple, practical rule for choosing just the right amount of undersmoothing. (See the end of this chapter for more discussion on this point.) We will take the first approach and content ourselves with finding a confidence band for $\bar{r}_{n}$.
5.98 Example. To understand the implications of estimating $\bar{r}_{n}$ instead of $r$, consider the following example. Let

$$
r(x)=\phi(x ; 2,1)+\phi(x ; 4,0.5)+\phi(x ; 6,0.1)+\phi(x ; 8,0.05)
$$

where $\phi(x ; m, s)$ denotes a Normal density function with mean $m$ and variance $s^{2}$. Figure 5.10 shows the true function (top left), a locally linear estimate $\widehat{r}_{n}$ (top right) based on 100 observations $Y_{i}=r(i / 10)+.2 N(0,1), i=1, \ldots, 100$, with bandwidth $h=0.27$, the function $\bar{r}_{n}(x)=\mathbb{E}\left(\widehat{r}_{n}(x)\right)$ (bottom left) and the difference $r(x)-\bar{r}_{n}(x)$ (bottom right). We see that $\bar{r}_{n}$ (dashed line) smooths out the peaks. Comparing the top right and bottom left plot, it is clear that $\widehat{r}_{n}(x)$ is actually estimating $\bar{r}_{n}$ not $r(x)$. Overall, $\bar{r}_{n}$ is quite similar to $r(x)$ except that $\bar{r}_{n}$ omits some of the fine details of $r$.

Constructing Confidence Bands. Assume that $\widehat{r}_{n}(x)$ is a linear smoother, so that $\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x)$. Then,

$$
\bar{r}(x)=\mathbb{E}\left(\widehat{r}_{n}(x)\right)=\sum_{i=1}^{n} \ell_{i}(x) r\left(x_{i}\right) .
$$

For now, assume that $\sigma^{2}(x)=\sigma^{2}=\mathbb{V}\left(\epsilon_{i}\right)$ is constant. Then,

$$
\mathbb{V}\left(\widehat{r}_{n}(x)\right)=\sigma^{2}\|\ell(x)\|^{2}
$$

[^7]

FIGURE 5.10. The true function (top left), an estimate $\widehat{r}_{n}$ (top right) based on 100 observations, the function $\bar{r}_{n}(x)=\mathbb{E}\left(\hat{r}_{n}(x)\right)$ (bottom left) and the difference $r(x)-\bar{r}_{n}(x)$ (bottom right).

We will consider a confidence band for $\bar{r}_{n}(x)$ of the form

$$
\begin{equation*}
I(x)=\left(\widehat{r}_{n}(x)-c \widehat{\sigma}\|\ell(x)\|, \widehat{r}_{n}(x)+c \widehat{\sigma}\|\ell(x)\|\right) \tag{5.99}
\end{equation*}
$$

for some $c>0$ and $a \leq x \leq b$.
We follow the approach in Sun and Loader (1994). First suppose that $\sigma$ is known. Then,

$$
\begin{aligned}
\mathbb{P}(\bar{r}(x) \notin I(x) \text { for some } x \in[a, b]) & =\mathbb{P}\left(\max _{x \in[a, b]} \frac{|\widehat{r}(x)-\bar{r}(x)|}{\sigma| | \ell(x) \|}>c\right) \\
=\mathbb{P}\left(\max _{x \in[a, b]} \frac{\left|\sum_{i} \epsilon_{i} \ell_{i}(x)\right|}{\sigma\|\ell(x)\|}>c\right) & =\mathbb{P}\left(\max _{x \in[a, b]}|W(x)|>c\right)
\end{aligned}
$$

where $W(x)=\sum_{i=1}^{n} Z_{i} T_{i}(x), Z_{i}=\epsilon_{i} / \sigma \sim N(0,1)$ and $T_{i}(x)=\ell_{i}(x) /\|\ell(x)\|$. Now, $W(x)$ is a Gaussian process. ${ }^{7}$ To find $c$, we need to be able to compute the distribution of the maximum of a Gaussian process. Fortunately, this is a well-studied problem. In particular, Sun and Loader (1994) showed that

$$
\begin{equation*}
\mathbb{P}\left(\max _{x}\left|\sum_{i=1}^{n} Z_{i} T_{i}(x)\right|>c\right) \approx 2(1-\Phi(c))+\frac{\kappa_{0}}{\pi} e^{-c^{2} / 2} \tag{5.100}
\end{equation*}
$$

for large $c$, where

$$
\begin{equation*}
\kappa_{0}=\int_{a}^{b}\left\|T^{\prime}(x)\right\| d x \tag{5.101}
\end{equation*}
$$

$T^{\prime}(x)=\left(T_{1}^{\prime}(x), \ldots, T_{n}^{\prime}(x)\right)$ and $T_{i}^{\prime}(x)=\partial T_{i}(x) / \partial x$. An approximation for $\kappa_{0}$ is given in Exercise 20. Equation (5.100) is called the tube formula. A derivation is outlined in the appendix. If we choose $c$ to solve

$$
\begin{equation*}
2(1-\Phi(c))+\frac{\kappa_{0}}{\pi} e^{-c^{2} / 2}=\alpha \tag{5.102}
\end{equation*}
$$

then we get the desired simultaneous confidence band. If $\sigma$ is unknown we use an estimate $\widehat{\sigma}$. Sun and Loader suggest replacing the right-hand side of (5.100) with

$$
\mathbb{P}\left(\left|T_{m}\right|>c\right)+\frac{\kappa_{0}}{\pi}\left(1+\frac{c^{2}}{m}\right)^{-m / 2}
$$

where $T_{m}$ has a t-distribution with $m=n-\operatorname{tr}(L)$ degrees of freedom. For large $n$, (5.100) remains an adequate approximation.

Now suppose that $\sigma(x)$ is a function of $x$. Then,

$$
\mathbb{V}\left(\widehat{r}_{n}(x)\right)=\sum_{i=1}^{n} \sigma^{2}\left(x_{i}\right) \ell_{i}^{2}(x)
$$

In this case we take

$$
\begin{equation*}
I(x)=\widehat{r}_{n}(x) \pm c s(x) \tag{5.103}
\end{equation*}
$$

where

$$
s(x)=\sqrt{\sum_{i=1}^{n} \hat{\sigma}^{2}\left(x_{i}\right) \ell_{i}^{2}(x)}
$$

$\widehat{\sigma}(x)$ is an estimate of $\sigma(x)$ and $c$ is the constant defined in (5.102). If $\widehat{\sigma}(x)$ varies slowly with $x$, then $\sigma\left(x_{i}\right) \approx \sigma(x)$ for those $i$ such that $\ell_{i}(x)$ is large and so

$$
s(x) \approx \widehat{\sigma}(x)\|\ell(x)\| .
$$

[^8]Thus, an approximate confidence band is

$$
\begin{equation*}
I(x)=\widehat{r}_{n}(x) \pm c \widehat{\sigma}(x)\|\ell(x)\| \tag{5.104}
\end{equation*}
$$

For more details on these methods, see Faraway and Sun (1995).
5.105 Example. Figure 5.11 shows simultaneous 95 percent confidence bands for the CMB data using a local linear fit. The bandwidth was chosen using cross-validation. We find that $\kappa_{0}=38.85$ and $c=3.33$. In the top plot, we assumed a constant variance when constructing the band. In the bottom plot, we did not assume a constant variance when constructing the band. We see that if we do not take into account the nonconstant variance, we overestimate the uncertainty for small $x$ and we underestimate the uncertainty for large $x$.
-



FIGURE 5.11. Local linear fit with simultaneous 95 percent confidence bands. The band in the top plot assumes constant variance $\sigma^{2}$. The band in the bottom plot allows for nonconstant variance $\widehat{\sigma}^{2}(x)$.
5.106 Remark. We have ignored the uncertainty introduced due to the choice of smoothing parameter. We can adjust for this extra uncertainty as follows. When choosing the smoothing parameter $h$, restrict the search to a finite set $\mathcal{H}_{n}$ with $m=m(n)$ points. Construct the confidence band at level $\alpha / m$. Thus, replace $\alpha$ with $\alpha / m$ on the right-hand side of equation (5.102). Then the Bonferroni inequality ${ }^{8}$ guarantees that the coverage is still at least $1-\alpha$.
5.107 Remark. There is a substantial literature on using bootstrap methods to get confidence bands. This requires a more sophisticated use of the bootstrap than that presented in Chapter 3. See Härdle and Marron (1991), Neumann and Polzehl (1998), Hall (1993), Faraway (1990), and Härdle and Bowman (1988), for example.

### 5.8 Average Coverage

One might argue that requiring bands to cover the function at all $x$ is too stringent. Wahba (1983), Nychka (1988) and Cummins et al. (2001) introduced a different type of coverage that we refer to as average coverage. Here we discuss a method for constructing average coverage bands based on an idea in Juditsky and Lambert-Lacroix (2003).

Suppose we are estimating $r(x)$ over the interval $[0,1]$. Define the average coverage of a band ( $\ell, u$ ) by

$$
C=\int_{0}^{1} \mathbb{P}(r(x) \in[\ell(x), u(x)]) d x
$$

In Chapters 7 and 8 we present methods for constructing confidence balls for $r$. These are sets $\mathcal{B}_{n}(\alpha)$ of the form

$$
\mathcal{B}_{n}(\alpha)=\left\{r:\left\|\widehat{r}_{n}-r\right\| \leq s_{n}(\alpha)\right\}
$$

such that

$$
\mathbb{P}\left(r \in \mathcal{B}_{n}(\alpha)\right) \geq 1-\alpha .
$$

Given such a confidence ball, let

$$
\begin{equation*}
\ell(x)=\widehat{r}_{n}(x)-s_{n}(\alpha / 2) \sqrt{\frac{2}{\alpha}}, \quad u(x)=\widehat{r}_{n}(x)+s_{n}(\alpha / 2) \sqrt{\frac{2}{\alpha}} . \tag{5.108}
\end{equation*}
$$

We now show that these bands have average coverage at least $1-\alpha$. First, note that $C=\mathbb{P}(r(U) \in[\ell(U), u(U)])$ where $U \sim \operatorname{Unif}(0,1)$ is independent of

[^9]the data. Let $A$ be the event that $r \in \mathcal{B}_{n}(\alpha / 2)$. On the event $A,\left\|\widehat{r}_{n}-r\right\| \leq$ $s_{n}(\alpha / 2)$. Writing $s_{n}$ for $s_{n}(\alpha / 2)$ we have,
\[

$$
\begin{aligned}
1-C & =\mathbb{P}(r(U) \notin[\ell(U), u(U)])=\mathbb{P}\left(\left|\widehat{r}_{n}(U)-r(U)\right|>s_{n} \sqrt{\frac{2}{\alpha}}\right) \\
& =\mathbb{P}\left(\left|\widehat{r}_{n}(U)-r(U)\right|>s_{n} \sqrt{\frac{2}{\alpha}}, A\right)+\mathbb{P}\left(\left|\widehat{r}_{n}(U)-r(U)\right|>s_{n} \sqrt{\frac{2}{\alpha}}, A^{c}\right) \\
& \leq \mathbb{P}\left(\left|\widehat{r}_{n}(U)-r(U)\right|>s_{n} \sqrt{\frac{2}{\alpha}}, A\right)+\mathbb{P}\left(A^{c}\right) \\
& \leq \frac{\mathbb{E} I_{A}\left|\widehat{r}_{n}(U)-r(U)\right|^{2}}{s_{n}^{2} \frac{2}{\alpha}}+\frac{\alpha}{2}=\frac{\mathbb{E} I_{A} \int_{0}^{1}\left|\widehat{r}_{n}(u)-r(u)\right|^{2} d u}{s_{n}^{2} \frac{2}{\alpha}}+\frac{\alpha}{2} \\
& =\frac{\mathbb{E} I_{A}| | \widehat{r}_{n}-\left.r\right|^{2}}{s_{n}^{2} \frac{2}{\alpha}}+\frac{\alpha}{2} \leq \frac{s_{n}^{2}}{s_{n}^{2} \frac{2}{\alpha}}+\frac{\alpha}{2} \leq \alpha .
\end{aligned}
$$
\]

### 5.9 Summary of Linear Smoothing

At this point we have covered many points related to linear smoothing methods. It seems like a good time so summarize the steps needed to construct the estimate $\widehat{r}_{n}$ and a confidence band.

## Summary of Linear Smoothing

1. Choose a smoothing method such as local polynomial, spline, etc.

This amounts to choosing the form of the weights $\ell(x)=$ $\left(\ell_{1}(x), \ldots, \ell_{n}(x)\right)^{T}$. A good default choice is local linear smoothing as described in Theorem 5.60.
2. Choose the bandwidth $h$ by cross-validation using (5.35).
3. Estimate the variance function $\widehat{\sigma}^{2}(x)$ as described in Section 5.6.
4. Find $\kappa_{0}$ from (5.101) and find $c$ from (5.102).
5. An approximate $1-\alpha$ confidence band for $\bar{r}_{n}=\mathbb{E}\left(\widehat{r}_{n}(x)\right)$ is

$$
\begin{equation*}
\widehat{r}_{n}(x) \pm c \widehat{\sigma}(x)\|\ell(x)\| . \tag{5.109}
\end{equation*}
$$

5.110 Example (LIDAR). Recall the LIDAR data from Example 4.5 and Example 5.59. We find that $\kappa_{0} \approx 30$ and $c \approx 3.25$. The resulting bands are
shown in the lower right plot. As expected, there is much greater uncertainty for larger values of the covariate.

### 5.10 Local Likelihood and Exponential Families

If $Y$ is not real valued or $\epsilon$ is not Gaussian, then the basic regression model we have been using might not be appropriate. For example, if $Y \in\{0,1\}$ then it seems natural to use a Bernoulli model. In this section we discuss nonparametric regression for more general models. Before proceeding, we should point out that the basic model often does work well even in cases where $Y$ is not real valued or $\epsilon$ is not Gaussian. This is because the asymptotic theory does not really depend on $\epsilon$ being Gaussian. Thus, at least for large samples, it is worth considering using the tools we have already developed for these cases.

Recall that $Y$ has an exponential family distribution, given $x$, if

$$
\begin{equation*}
f(y \mid x)=\exp \left\{\frac{y \theta(x)-b(\theta(x))}{a(\phi)}+c(y, \phi)\right\} \tag{5.111}
\end{equation*}
$$

for some functions $a(\cdot), b(\cdot)$ and $c(\cdot, \cdot)$. Here, $\theta(\cdot)$ is called the canonical parameter and $\phi$ is called the dispersion parameter. It then follows that

$$
\begin{aligned}
r(x) & \equiv \mathbb{E}(Y \mid X=x)=b^{\prime}(\theta(x)) \\
\sigma^{2}(x) & \equiv \mathbb{V}(Y \mid X=x)=a(\phi) b^{\prime \prime}(\theta(x))
\end{aligned}
$$

The usual parametric form of this model is

$$
g(r(x))=x^{T} \beta
$$

for some known function $g$ called the link function. The model

$$
\begin{equation*}
Y \mid X=x \sim f(y \mid x), \quad g(\mathbb{E}(Y \mid X=x))=x^{T} \beta \tag{5.112}
\end{equation*}
$$

is called a generalized linear model.
For example, if $Y$ given $X=x$ is Binomial $(m, r(x))$ then

$$
f(y \mid x)=\binom{m}{y} r(x)^{y}(1-r(x))^{m-y}
$$

which has the form (5.111) with

$$
\theta(x)=\log \frac{r(x)}{1-r(x)}, \quad b(\theta)=m \log \left(1+e^{\theta}\right)
$$

and $a(\phi) \equiv 1$. Taking $g(t)=\log (t /(m-t))$ yields the logistic regression model. The parameters $\beta$ are usually estimated by maximum likelihood.

Let's consider a nonparametric version of logistic regression. For simplicity, we focus on local linear estimation. The data are $\left(x_{1}, Y_{1}\right), \ldots,\left(x_{n}, Y_{n}\right)$ where $Y_{i} \in\{0,1\}$. We assume that

$$
Y_{i} \sim \operatorname{Bernoulli}\left(r\left(x_{i}\right)\right)
$$

for some smooth function $r(x)$ for which $0 \leq r(x) \leq 1$. Thus, $\mathbb{P}\left(Y_{i}=1 \mid X_{i}=\right.$ $\left.x_{i}\right)=r\left(x_{i}\right)$ and $\mathbb{P}\left(Y_{i}=0 \mid X_{i}=x_{i}\right)=1-r\left(x_{i}\right)$. The likelihood function is

$$
\prod_{i=1}^{n} r\left(x_{i}\right)^{Y_{i}}\left(1-r\left(x_{i}\right)\right)^{1-Y_{i}}
$$

so, with $\xi(x)=\log (r(x) /(1-r(x)))$, the log-likelihood is

$$
\begin{equation*}
\ell(r)=\sum_{i=1}^{n} \ell\left(Y_{i}, \xi\left(x_{i}\right)\right) \tag{5.113}
\end{equation*}
$$

where

$$
\begin{align*}
\ell(y, \xi) & =\log \left[\left(\frac{e^{\xi}}{1+e^{\xi}}\right)^{y}\left(\frac{1}{1+e^{\xi}}\right)^{1-y}\right] \\
& =y \xi-\log \left(1+e^{\xi}\right) \tag{5.114}
\end{align*}
$$

To estimate the regression function at $x$ we approximate the regression function $r(u)$ for $u$ near $x$ by the local logistic function

$$
r(u) \approx \frac{e^{a_{0}+a_{1}(u-x)}}{1+e^{a_{0}+a_{1}(u-x)}}
$$

(compare with (5.15)). Equivalently, we approximate $\log (r(u) /(1-r(u)))$ with $a_{0}+a_{1}(x-u)$. Now define the local log-likelihood

$$
\begin{aligned}
\ell_{x}(a) & =\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right) \ell\left(Y_{i}, a_{0}+a_{1}\left(X_{i}-x\right)\right) \\
& =\sum_{i=1}^{n} K\left(\frac{x-X_{i}}{h}\right)\left(Y_{i}\left(a_{0}+a_{1}\left(X_{i}-x\right)\right)-\log \left(1+e^{a_{0}+a_{1}\left(X_{i}-x\right)}\right)\right)
\end{aligned}
$$

Let $\widehat{a}(x)=\left(\widehat{a}_{0}(x), \widehat{a}_{1}(x)\right)$ maximize $\ell_{x}$ which can be found by any convenient optimization routine such as Newton-Raphson. The nonparametric estimate of $r(x)$ is

$$
\begin{equation*}
\widehat{r}_{n}(x)=\frac{e^{\widehat{a}_{0}(x)}}{1+e^{\widehat{a}_{0}(x)}} \tag{5.115}
\end{equation*}
$$

The bandwidth can be chosen by using the leave-one-out log-likelihood cross-validation

$$
\begin{equation*}
\mathrm{CV}=\sum_{i=1}^{n} \ell\left(Y_{i}, \widehat{\xi}_{(-i)}\left(x_{i}\right)\right) \tag{5.116}
\end{equation*}
$$

where $\widehat{\xi}_{(-i)}(x)$ is the estimator obtained by leaving out $\left(x_{i}, Y_{i}\right)$. Unfortunately, there is no identity as in Theorem 5.34. There is, however, the following approximation from Loader (1999a). Recall the definition of $\ell(x, \xi)$ from (5.114) and let $\dot{\ell}(y, \xi)$ and $\ddot{\ell}(y, \xi)$ denote first and second derivatives of $\ell(y, \xi)$ with respect to $\xi$. Thus,

$$
\begin{aligned}
\dot{\ell}(y, \xi) & =y-p(\xi) \\
\ddot{\ell}(y, \xi) & =-p(\xi)(1-p(\xi))
\end{aligned}
$$

where $p(\xi)=e^{\xi} /\left(1+e^{\xi}\right)$. Define matrices $X_{x}$ and $W_{x}$ as in (5.54) and let $V_{x}$ be a diagonal matrix with $j^{\text {th }}$ diagonal entry equal to $-\ddot{\ell}\left(Y_{i}, \widehat{a}_{0}+\widehat{a}_{1}\left(x_{j}-x_{i}\right)\right)$. Then,

$$
\begin{equation*}
\mathrm{CV} \approx \ell_{x}(\widehat{a})+\sum_{i=1}^{n} m\left(x_{i}\right)\left(\dot{\ell}\left(Y_{i}, \widehat{a}_{0}\right)\right)^{2} \tag{5.117}
\end{equation*}
$$

where

$$
\begin{equation*}
m(x)=K(0) e_{1}^{T}\left(X_{x}^{T} W_{x} V_{x} X_{x}\right)^{-1} e_{1} \tag{5.118}
\end{equation*}
$$

and $e_{1}=(1,0, \ldots, 0)^{T}$. The effective degrees of freedom is

$$
\nu=\sum_{i=1}^{n} m\left(x_{i}\right) \mathbb{E}\left(-\ddot{\ell}\left(Y_{i}, \widehat{a}_{0}\right)\right) .
$$

5.119 Example. Figure 5.12 shows the local linear logistic regression estimator for an example generated from the model $Y_{i} \sim \operatorname{Bernoulli}\left(r\left(x_{i}\right)\right)$ with $r(x)=e^{3 \sin (x)} /\left(1+e^{3 \sin (x)}\right)$. The solid line is the true function $r(x)$. The dashed line is the local linear logistic regression estimator. We also computed the local linear regression estimator which ignores the fact that the data are Bernoulli. The dotted line is the resulting local linear regression estimator. ${ }^{9}$ cross-validation was used to select the bandwidth in both cases.
5.120 Example. We introduced the BPD data in Example 4.6. The outcome $Y$ is presence or absence of BPD and the covariate is $x=$ birth weight. The estimated logistic regression function (solid line) $r\left(x ; \widehat{\beta}_{0}, \widehat{\beta}_{1}\right)$ together with the

[^10]

FIGURE 5.12. Local linear logistic regression. The solid line is the true regression function $r(x)=\mathbb{P}(Y=1 \mid X=x)$. The dashed line is the local linear logistic regression estimator. The dotted line is the local linear regression estimator.
data are shown in Figure 5.13. Also shown are two nonparametric estimates. The dashed line is the local likelihood estimator. The dotted line is the local linear estimator which ignores the binary nature of the $Y_{i}$ 's. Again we see that there is not a dramatic difference between the local logistic model and the local linear model.

### 5.11 Scale-Space Smoothing

There is another approach to smoothing championed by Chaudhuri and Marron (1999) and Chaudhuri and Marron (2000) called scale-space smoothing that eschews the idea of selecting a single bandwidth. Let $\widehat{r}_{h}(x)$ denote an estimator using bandwidth $h$. The idea is to regard $\widehat{r}_{h}(x)$ as an estimator of $r_{h}(x) \equiv \mathbb{E}\left(\widehat{r}_{h}(x)\right)$, as we did in Section 5.7. But rather than choosing a single bandwidth, we examine $\widehat{r}_{h}$ over a set of bandwidths $h$ as a way of exploring the scale-space surface

$$
\mathcal{S}=\left\{r_{h}(x), x \in \mathcal{X}, h \in \mathcal{H}\right\}
$$

where $\mathcal{X}$ is the range of $x$ and $\mathcal{H}$ is the range of $h$.
One way to summarize the estimated scale-space surface

$$
\widehat{\mathcal{S}}=\left\{\widehat{r}_{h}(x), x \in \mathcal{X}, h \in \mathcal{H}\right\}
$$

is to isolate important shape summaries. For example, Chaudhuri and Marron (1999) look for points $x$ where $r_{h}^{\prime}(x)=0$ by using $\widehat{r}_{h}^{\prime}(x)$ as a set of test


FIGURE 5.13. The BPD data. The data are shown with small vertical lines. The estimates are from logistic regression (solid line), local likelihood (dashed line) and local linear regression (dotted line).
statistics. They call the resulting method SiZer (significant zero crossings of derivatives).

### 5.12 Multiple Regression

Suppose now that the covariate is $d$-dimensional,

$$
x_{i}=\left(x_{i 1}, \ldots, x_{i d}\right)^{T}
$$

The regression equation takes the form

$$
\begin{equation*}
Y=r\left(x_{1}, \ldots, x_{d}\right)+\epsilon \tag{5.121}
\end{equation*}
$$

In principle, all the methods we have discussed carry over to this case easily. Unfortunately, the risk of a nonparametric regression estimator increases rapidly with the dimension $d$. This is the curse of dimensionality that we discussed in Section 4.5. It is worth revisiting this point now. In a one-dimensional
problem, the optimal rate of convergence of a nonparametric estimator is $n^{-4 / 5}$ if $r$ is assumed to have an integrable second derivative. In $d$ dimensions the optimal rate of convergence is is $n^{-4 /(4+d)}$. Thus, the sample size $m$ required for a $d$-dimensional problem to have the same accuracy as a sample size $n$ in a one-dimensional problem is $m \propto n^{c d}$ where $c=(4+d) /(5 d)>0$. This implies the following fact:

## To maintain a given degree of accuracy of an estimator, the sample size must increase exponentially with the dimension $d$.

Put another way, confidence bands get very large as the dimension $d$ increases. Nevertheless, let us press on and see how we might estimate the regression function.

Local Regression. Consider local linear regression. The kernel function $K$ is now a function of $d$ variables. Given a nonsingular positive definite $d \times d$ bandwidth matrix $H$, we define

$$
K_{H}(x)=\frac{1}{|H|^{1 / 2}} K\left(H^{-1 / 2} x\right)
$$

Often, one scales each covariate to have the same mean and variance and then we use the kernel

$$
h^{-d} K(\|x\| / h)
$$

where $K$ is any one-dimensional kernel. Then there is a single bandwidth parameter $h$. At a target value $x=\left(x_{1}, \ldots, x_{d}\right)^{T}$, the local sum of squares is given by

$$
\begin{equation*}
\sum_{i=1}^{n} w_{i}(x)\left(Y_{i}-a_{0}-\sum_{j=1}^{d} a_{j}\left(x_{i j}-x_{j}\right)\right)^{2} \tag{5.122}
\end{equation*}
$$

where

$$
w_{i}(x)=K\left(\left\|x_{i}-x\right\| / h\right) .
$$

The estimator is

$$
\begin{equation*}
\widehat{r}_{n}(x)=\widehat{a}_{0} \tag{5.123}
\end{equation*}
$$

where $\widehat{a}=\left(\widehat{a}_{0}, \ldots, \widehat{a}_{d}\right)^{T}$ is the value of $a=\left(a_{0}, \ldots, a_{d}\right)^{T}$ that minimizes the weighted sums of squares. The solution $\widehat{a}$ is

$$
\begin{equation*}
\widehat{a}=\left(X_{x}^{T} W_{x} X_{x}\right)^{-1} X_{x}^{T} W_{x} Y \tag{5.124}
\end{equation*}
$$

where

$$
X_{x}=\left(\begin{array}{cccc}
1 & x_{11}-x_{1} & \cdots & x_{1 d}-x_{d} \\
1 & x_{21}-x_{1} & \cdots & x_{2 d}-x_{d} \\
\vdots & \vdots & \ddots & \vdots \\
1 & x_{n 1}-x_{1} & \cdots & x_{n d}-x_{d}
\end{array}\right)
$$

and $W_{x}$ is the diagonal matrix whose $(i, i)$ element is $w_{i}(x)$.
The theoretical properties of local polynomial regression in higher dimensions is discussed in Ruppert and Wand (1994). The main result is as follows.
5.125 Theorem (Ruppert and Wand, 1994). Let $\widehat{r}_{n}$ be the multivariate local linear estimator with bandwidth matrix $H$ and assume the regularity conditions given in Ruppert and Wand (1994). Suppose that $x$ is a nonboundary point. Conditional on $X_{1}, \ldots, X_{n}$ we have the following: The bias of $\widehat{r}_{n}(x)$ is

$$
\begin{equation*}
\frac{1}{2} \mu_{2}(K) \operatorname{trace}(H \mathcal{H})+o_{P}(\operatorname{trace}(H)) \tag{5.126}
\end{equation*}
$$

where $\mathcal{H}$ is the matrix of second partial derivatives of $r$ evaluated at $x$ and $\mu_{2}(K)$ is the scalar defined by the equation $\int u u^{T} K(u) d u=\mu_{2}(K) I$. The variance of $\widehat{r}_{n}(x)$ is

$$
\begin{equation*}
\frac{\sigma^{2}(x) \int K(u)^{2} d u}{n|H|^{1 / 2} f(x)}\left(1+o_{P}(1)\right) . \tag{5.127}
\end{equation*}
$$

Also, the bias at the boundary is the same order as in the interior, namely, $O_{P}(\operatorname{trace}(H))$.

Thus we see that in higher dimensions, local linear regression still avoids excessive boundary bias and design bias.

Splines. If we take a spline approach, we need to define splines in higher dimensions. For $d=2$ we minimize

$$
\sum_{i}\left(Y_{i}-\widehat{r}_{n}\left(x_{i 1}, x_{i 2}\right)\right)^{2}+\lambda J(r)
$$

where

$$
J(r)=\iint\left[\left(\frac{\partial^{2} r(x)}{\partial x_{1}^{2}}\right)+2\left(\frac{\partial^{2} r(x)}{\partial x_{1} \partial x_{2}}\right)+\left(\frac{\partial^{2} r(x)}{\partial x_{2}^{2}}\right)\right] d x_{1} d x_{2}
$$

The minimizer $\widehat{r}_{n}$ is called a thin-plate spline. It is hard to describe and even harder (but certainly not impossible) to fit. See Green and Silverman (1994) for more details.

Additive Models. Interpreting and visualizing a high-dimensional fit is difficult. As the number of covariates increases, the computational burden becomes prohibitive. Sometimes, a more fruitful approach is to use an additive model. An additive model is a model of the form

$$
\begin{equation*}
Y=\alpha+\sum_{j=1}^{d} r_{j}\left(x_{j}\right)+\epsilon \tag{5.128}
\end{equation*}
$$

where $r_{1}, \ldots, r_{d}$ are smooth functions. The model (5.128) is not identifiable since we can add any constant to $\alpha$ and subtract the same constant from one of the $r_{j}$ 's without changing the regression function. This problem can be fixed in a number of ways, perhaps the easiest being to set $\widehat{\alpha}=\bar{Y}$ and then regard the $r_{j}$ 's as deviations from $\bar{Y}$. In this case we require that $\sum_{i=1}^{n} \widehat{r}_{j}\left(x_{i}\right)=0$ for each $j$.

The additive model is clearly not as general as fitting $r\left(x_{1}, \ldots, x_{d}\right)$ but it is much simpler to compute and to interpret and so it is often a good starting point. This is a simple algorithm for turning any one-dimensional regression smoother into a method for fitting additive models. It is called backfitting.

The Backfitting Algorithm
Initialization: set $\widehat{\alpha}=\bar{Y}$ and set initial guesses for $\widehat{r}_{1}, \ldots, \widehat{r}_{d}$.
Iterate until convergence: for $j=1, \ldots, d$ :

- Compute $\widetilde{Y}_{i}=Y_{i}-\widehat{\alpha}-\sum_{k \neq j} r_{k}\left(x_{i}\right), i=1, \ldots, n$.
- Apply a smoother to $\widetilde{Y}_{i}$ on $x_{j}$ to obtain $\widehat{r}_{j}$.
- Set $\widehat{r}_{j}(x)$ equal to $\widehat{r}_{j}(x)-n^{-1} \sum_{i=1}^{n} \widehat{r}_{j}\left(x_{i}\right)$.
5.129 Example. Let us revisit Example 4.7 involving three covariates and one response variable. The data are plotted in Figure 5.14. Recall that the data are 48 rock samples from a petroleum reservoir, the response is permeability (in milli-Darcies) and the covariates are: the area of pores (in pixels out of 256 by 256), perimeter in pixels and shape (perimeter $/ \sqrt{\text { area). }}$. The goal is to predict permeability from the three covariates. First we fit the additive model

$$
\text { permeability }=r_{1}(\text { area })+r_{2}(\text { perimeter })+r_{3}(\text { shape })+\epsilon .
$$

We could scale each covariate to have the same variance and then use a common bandwidth for each covariate. Instead, I took the more adventurous approach of performing cross-validation to choose a bandwidth $h_{j}$ for covariate
$x_{j}$ during each iteration of backfitting. I am not aware of any theory that guarantees convergence if the smoothing parameters are changed this way during the algorithm. Nonetheless, the bandwidths and the functions estimates converged rapidly. The estimates of $r_{1}, r_{2}$ and $r_{3}$ are shown in Figure $5.15 . \bar{Y}$ was added to each function before plotting it. Next consider a three-dimensional local linear fit (5.123). After scaling each covariate to have mean 0 and variance 1 , we found that the bandwidth $h \approx 3.2$ minimized the cross-validation score. The residuals from the additive model and the full three-dimensional local linear fit are shown in Figure 5.16. Apparently, the fitted values are quite similar suggesting that the generalized additive model is adequate.


FIGURE 5.14. The rock data.

Projection Pursuit. Friedman and Stuetzle (1981) introduced another method for dealing with high-dimensional regression called projection pursuit regression. The idea is to approximate the regression function


FIGURE 5.15. The rock data. The plots show $\widehat{r}_{1}, \widehat{r}_{2}$, and $\widehat{r}_{3}$ for the additive model $Y=\widehat{r}_{1}\left(x_{1}\right)+\widehat{r}_{2}\left(x_{2}\right)+\widehat{r}_{3}\left(x_{3}\right)+\epsilon$.
$r\left(x_{1}, \ldots, x_{p}\right)$ with a function of the form

$$
\mu+\sum_{m=1}^{M} r_{m}\left(z_{m}\right)
$$

where

$$
z_{m}=\alpha_{m}^{T} x
$$

and each $\alpha_{m}$ is a unit vector (length one) for $m=1, \ldots, M$. Note that each $z_{m}$ is the projection of $x$ into a subspace. The direction vector $\alpha$ is chosen at each stage to minimize the fraction of unexplained variance. In more detail, let $S(\cdot)$ denote the mapping that outputs $n$ fitted values from some smoothing method, given the $Y_{i} \mathrm{~s}$ and some one-dimensional covariate values $z_{1}, \ldots, z_{n}$. Let $\widehat{\mu}=\bar{Y}$ and replace $Y_{i}$ with $Y_{i}-\bar{Y}$. Hence, the $Y_{i}$ 's now have mean 0 . Similarly, scale the covariates so that they each have the same variance. Then do the following steps:


FIGURE 5.16. The residuals for the rock data. Top left: residuals from the additive model. Top right: qq-plot of the residuals from the additive model. Bottom left: residuals from the multivariate local linear model. Bottom right: residuals from the two fits plotted against each other.

Step 1: Initialize residuals $\widehat{\epsilon}_{i}=Y_{i}, i=1, \ldots, n$ and set $m=0$.
Step 2: Find the direction (unit vector) $\alpha$ that maximizes

$$
I(\alpha)=1-\frac{\sum_{i=1}^{n}\left(\widehat{\epsilon}_{i}-S\left(\alpha^{T} x_{i}\right)\right)^{2}}{\sum_{i=1}^{n} \widehat{\epsilon}_{i}^{2}}
$$

and set $z_{m i}=\alpha^{T} x_{i}, \widehat{r}_{m}\left(z_{m i}\right)=S\left(z_{m i}\right)$.
Step 3: Set $m=m+1$ and update the residuals:

$$
\widehat{\epsilon}_{i} \longleftarrow \widehat{\epsilon}_{i}-\widehat{r}_{m}\left(z_{m i}\right)
$$

If $m=M$ stop, otherwise go back to Step 2.
5.130 Example. If we apply projection pursuit regression to the rock data with $M=3$ we get the functions $\widehat{r}_{1}, \widehat{r}_{2}, \widehat{r}_{3}$ shown in 5.17. The fitting was done
using the ppr command in R and each fit was obtained using smoothing splines where the smoothing parameter is chosen by generalized cross-validation. The direction vectors are

$$
\alpha_{1}=(.99, .07, .08)^{T}, \alpha_{2}=(.43, .35, .83)^{T}, \alpha_{3}=(.74,-.28,-.61)^{T}
$$

Thus, $z_{1}=.99$ area +.07 peri +.08 shape and so on. If we keep adding terms to the model, the residual sums of squares will keep getting smaller. The bottom left plot in Figure 5.17 shows the residual sums of squares as a function of the number of terms $M$. We see that after including one or two terms in the model, further terms add little. We could try to choose an optimal $M$ by using cross-validation.


FIGURE 5.17. Projection pursuit applied to the rock data. The plots show $\widehat{r}_{1}, \widehat{r}_{2}$, and $\widehat{r}_{3}$.

Regression Trees. A regression tree is a model of the form

$$
\begin{equation*}
r(x)=\sum_{m=1}^{M} c_{m} I\left(x \in R_{m}\right) \tag{5.131}
\end{equation*}
$$

where $c_{1}, \ldots, c_{M}$ are constants and $R_{1}, \ldots, R_{M}$ are disjoint rectangles that partition the space of covariates. Tree models were introduced by Morgan and Sonquist (1963) and Breiman et al. (1984). The model is fitted in a recursive manner that can be represented as a tree; hence the name. Our description follows Section 9.2 of Hastie et al. (2001).

Denote a generic covariate value by $x=\left(x_{1}, \ldots, x_{j}, \ldots, x_{d}\right)$. The covariate for the $i^{\text {th }}$ observation is $x_{i}=\left(x_{i 1}, \ldots, x_{i j}, \ldots, x_{i d}\right)$. Given a covariate $j$ and a split point $s$ we define the rectangles $R_{1}=R_{1}(j, s)=\left\{x: x_{j} \leq s\right\}$ and $R_{2}=R_{2}(j, s)=\left\{x: x_{j}>s\right\}$ where, in this expression, $x_{j}$ refers the the $j^{\text {th }}$ covariate not the $j^{\text {th }}$ observation. Then we take $c_{1}$ to be the average of all the $Y_{i}$ 's such that $x_{i} \in R_{1}$ and $c_{2}$ to be the average of all the $Y_{i}$ 's such that $x_{i} \in R_{2}$. Notice that $c_{1}$ and $c_{2}$ minimize the sums of squares $\sum_{x_{i} \in R_{1}}\left(Y_{i}-c_{1}\right)^{2}$ and $\sum_{x_{i} \in R_{2}}\left(Y_{i}-c_{2}\right)^{2}$. The choice of which covariate $x_{j}$ to split on and which split point $s$ to use is based on minimizing the residual sums if squares. The splitting process is on repeated on each rectangle $R_{1}$ and $R_{2}$.

Figure 5.18 shows a simple example of a regression tree; also shown are the corresponding rectangles. The function estimate $\widehat{r}$ is constant over the rectangles.

Generally one grows a very large tree, then the tree is pruned to form a subtree by collapsing regions together. The size of the tree is a tuning parameter chosen as follows. Let $N_{m}$ denote the number of points in a rectangle $R_{m}$ of a subtree $T$ and define

$$
c_{m}=\frac{1}{N_{m}} \sum_{x_{i} \in R_{m}} Y_{i}, \quad Q_{m}(T)=\frac{1}{N_{m}} \sum_{x_{i} \in R_{m}}\left(Y_{i}-c_{m}\right)^{2} .
$$

Define the complexity of $T$ by

$$
\begin{equation*}
C_{\alpha}(T)=\sum_{m=1}^{|T|} N_{m} Q_{m}(T)+\alpha|T| \tag{5.132}
\end{equation*}
$$

where $\alpha>0$ and $|T|$ is the number of terminal nodes of the tree. Let $T_{\alpha}$ be the smallest subtree that minimizes $C_{\alpha}$. The value $\widehat{\alpha}$ of $\alpha$ can be chosen by cross-validation. The final estimate is based on the tree $T_{\widehat{\alpha}}$.
5.133 Example. Figure 5.19 shows a tree for the rock data. Notice that the variable shape does not appear in the tree. This means that the shape variable



FIGURE 5.18. A regression tree for two covariates $x_{1}$ and $x_{2}$. The function estimate is $\widehat{r}(x)=c_{1} I\left(x \in R_{1}\right)+c_{2} I\left(x \in R_{2}\right)+c_{3} I\left(x \in R_{3}\right)$ where $R_{1}, R_{2}$ and $R_{3}$ are the rectangles shown in the lower plot.
was never the optimal covariate to split on in the algorithm. The result is that tree only depends on area and peri. This illustrates an important feature of tree regression: it automatically performs variable selection in the sense that a covariate $x_{j}$ will not appear in the tree if the algorithm finds that the variable is not important.


FIGURE 5.19. Regression tree for the rock data.

MARS. Regression trees are discontinuous and they do not easily fit main effects. (In contrast, additive models fit only main effects unless interactions are put in the model.) MARS - which stands for multivariate adaptive regression splines - was introduced by Friedman (1991) as an attempt to improve on regression trees.

The MARS algorithm is as follows. We follow Section 9.4 of Hastie et al. (2001). Define

$$
\ell(x, j, t)=\left(x_{j}-t\right) I\left(x_{j}>t\right), \quad r(x, j, t)=\left(t-x_{j}\right) I\left(x_{j}<t\right) .
$$

Both $\ell(x, j, t)$ and $r(x, j, t)$ are functions of the whole vector $x=\left(x_{1}, \ldots, x_{d}\right)$ but their values only depend on the $j^{\text {th }}$ component $x_{j}$. Let

$$
\mathcal{C}_{j}=\left\{\ell(x, j, t), r(x, j, t), t \in\left\{x_{1 j}, \ldots, x_{n j}\right\}\right\} .
$$

Thus, $\mathcal{C}_{j}$ is a set of linear splines that depend only on $x_{j}$, with knots at the observations. Now let $\mathcal{C}=\bigcup_{j=1}^{d} \mathcal{C}_{j}$. A MARS model is of the form

$$
\begin{equation*}
r(x)=\beta_{0}+\sum_{m=1}^{M} \beta_{m} h_{m}(x) \tag{5.134}
\end{equation*}
$$

where every function $h_{m}$ is either in $\mathcal{C}$ or is the product of two or more such functions. The model is fitted in a forward, stepwise manner, much like regression trees. See Hastie et al. (2001) for more details.

Tensor Product Models. Another class of models for multiple regression are of the form

$$
\begin{equation*}
r(x)=\sum_{m=1}^{M} \beta_{m} h_{m}(x) \tag{5.135}
\end{equation*}
$$

where each $h_{m}$ is a basis function in a tensor product space. These models will be considered in Chapter 8.

### 5.13 Other Issues

Here we discuss a few other issues related to nonparametric regression.

Plug-In Bandwidths. An alternative method to cross-validation for choosing the bandwidth is to use a plug-in bandwidth. The idea is to write down a formula for the asymptotically optimal bandwidth and then insert estimates of the unknown quantities in the formula. We describe one possible approach based on Section 4.2 of Fan and Gijbels (1996).

When using local linear regression, and assuming that the $X_{i}$ 's are randomly selected from some density $f(x)$, the (asymptotically) optimal bandwidth is

$$
\begin{equation*}
h_{*}=\left(\frac{C \int \frac{\sigma^{2}(x)}{f(x)} d x}{n \int\left(r^{(2)}(x)\right)^{2} d x}\right)^{1 / 5} \tag{5.136}
\end{equation*}
$$

where

$$
\begin{equation*}
C=\frac{\int K^{2}(t) d t}{\left(\int K(t) d t\right)^{2}} \tag{5.137}
\end{equation*}
$$

and $r^{(2)}$ is the second derivative of $r$. To get a crude estimate of $h_{*}$ proceed as follows. Fit a global quartic

$$
\widetilde{r}(x)=\widehat{\beta}_{0}+\widehat{\beta}_{1} x+\widehat{\beta}_{2} x^{2}+\widehat{\beta}_{3} x^{3}+\widehat{\beta}_{4} x^{4}
$$

using least squares and let $\widetilde{\sigma}^{2}=n^{-1} \sum_{i=1}^{n}\left(Y_{i}-\widetilde{r}\left(x_{i}\right)\right)^{2}$. We call $\widetilde{r}(x)$ the pilot estimate. Let $(a, b)$ denote the range of the $X_{i}$ 's and approximate $f$ with a uniform over $(a, b)$. Then,

$$
\begin{aligned}
n \int_{a}^{b} r^{(2)}(x)^{2} d x & =n \int_{a}^{b} \frac{r^{(2)}(x)^{2}}{f(x)} f(x) d x \\
& \approx \sum_{i=1}^{n} \frac{r^{(2)}\left(X_{i}\right)^{2}}{f\left(X_{i}\right)}=(b-a) \sum_{i=1}^{n}\left(r^{(2)}\left(X_{i}\right)\right)^{2}
\end{aligned}
$$

and we estimate $h_{*}$ by

$$
\begin{equation*}
h_{*} \approx\left(\frac{C \widetilde{\sigma}^{2}(b-a)}{\sum_{i=1}^{n}\left(\widetilde{r}^{(2)}\left(X_{i}\right)\right)^{2}}\right)^{1 / 5} \tag{5.138}
\end{equation*}
$$

Testing for Linearity. A nonparametric estimator $\widehat{r}_{n}$ can be used to construct a test to see whether a linear fit is adequate. Consider testing

$$
H_{0}: r(x)=\beta_{0}+\beta_{1} x \text { for some } \beta_{0}, \beta_{1}
$$

versus the alternative that $H_{0}$ is false.
Denote the hat matrix from fitting the linear model by $H$ and the smoothing matrix from fitting the nonparametric regression by $L$. Let

$$
T=\frac{\|L Y-H Y\| / \lambda}{\widehat{\sigma}^{2}}
$$

where $\lambda=\operatorname{tr}\left((L-H)^{T}(L-H)\right)$ and $\widehat{\sigma}^{2}$ is defined by (5.86). Loader (1999a) points out that under $H_{0}$, the F-distribution with $\nu$ and $n-2 \nu_{1}+\nu_{2}$ degrees of freedom provides a rough approximation to the distribution of $T$. Thus we would reject $H_{0}$ at level $\alpha$ if $T>F_{\nu, n-2 \nu_{1}+\nu_{2}, \alpha}$. A more rigorous test that uses bootstrapping to estimate the null distribution is described in Härdle and Mammen (1993).

As with any test, the failure to reject $H_{0}$ should not be regarded as proof that $H_{0}$ is true. Rather, it indicates that the data are not powerful enough to detect deviations from $H_{0}$. In such cases, a linear fit might be considered a reasonable tentative model. Of course, making such decisions based solely on the basis of a test can be dangerous.

Optimality. Local linear estimators have certain optimality properties. We highlight a few results from Fan and Gijbels (1996). Let $x_{0}$ be an interior
(nonboundary) point and let

$$
\mathcal{F}=\left\{r:\left|r(x)-r\left(x_{0}\right)-\left(x-x_{0}\right) r^{\prime}\left(x_{0}\right)\right| \leq C\left|x-x_{0}\right|\right\} .
$$

Suppose that the covariates $X$ is random with density $f$ that is positive at $x_{0}$. Also, assume that the variance function $\sigma(x)$ is continuous at $x_{0}$. Let $\mathcal{L}$ denote all linear estimators of $r\left(x_{0}\right)$. The linear minimax risk is defined by

$$
\begin{equation*}
R_{n}^{L}=\inf _{\hat{\theta} \in \mathcal{L}} \sup _{r \in \mathcal{F}} \mathbb{E}\left(\left(\widehat{\theta}-r\left(x_{0}\right)\right)^{2} \mid X_{1}, \ldots, X_{n}\right) \tag{5.139}
\end{equation*}
$$

Fan and Gijbels (1996) show that

$$
\begin{equation*}
R_{n}^{L}=\frac{3}{4} 15^{-1 / 5}\left(\frac{\sqrt{C} \sigma^{2}\left(x_{0}\right)}{n f\left(x_{0}\right)}\right)^{4 / 5}\left(1+o_{P}(1)\right) \tag{5.140}
\end{equation*}
$$

Moreover, this risk is achieved by the local linear estimator $\widehat{r}_{*}$ with Epanechnikov kernel and bandwidth

$$
h_{*}=\left(\frac{15 \sigma^{2}\left(x_{0}\right)}{f\left(x_{0}\right) C^{2} n}\right)^{1 / 5}
$$

The minimax risk is defined by

$$
\begin{equation*}
R_{n}=\inf _{\hat{\theta}} \sup _{r \in \mathcal{F}} \mathbb{E}\left(\left(\widehat{\theta}-r\left(x_{0}\right)\right)^{2} \mid X_{1}, \ldots, X_{n}\right) \tag{5.141}
\end{equation*}
$$

where the infimum is now over all estimators. Fan and Gijbels (1996) show that $\widehat{r}_{*}$ is nearly minimax in the sense that

$$
\begin{equation*}
\frac{R_{n}}{\sup _{r \in \mathcal{F}} \mathbb{E}\left(\left(\widehat{r}^{*}\left(x_{0}\right)-r\left(x_{0}\right)\right)^{2} \mid X_{1}, \ldots, X_{n}\right)} \geq(0.894)^{2}+o_{P}(1) \tag{5.142}
\end{equation*}
$$

See Chapter 7 for more discussion on minimaxity.

Derivative Estimation. Suppose we want to estimate the $k^{\text {th }}$ derivative $r^{(k)}(x)$ of $r(x)$. Recall that the local polynomial estimator begins with the approximation

$$
r(u) \approx a_{0}+a_{1}(u-x)+\frac{a_{2}}{2}(u-x)^{2}+\cdots+\frac{a_{p}}{p!}(u-x)^{p} .
$$

Thus, $r^{(k)}(x) \approx a_{k}$ and we can estimate this with

$$
\begin{equation*}
\widehat{r}_{n}^{(k)}(x)=\widehat{a}_{k}=\sum_{i=1}^{n} \ell_{i}(x, k) Y_{i} \tag{5.143}
\end{equation*}
$$

where $\ell(x, k)^{T}=\left(\ell_{1}(x, k), \ldots, \ell_{n}(x, k)\right)$,

$$
\begin{aligned}
\ell(x, k)^{T} & =e_{k+1}^{T}\left(X_{x}^{T} W_{x} X_{x}\right)^{-1} X_{x}^{T} W_{x}, \\
e_{k+1} & =(\underbrace{0, \ldots, 0}_{k}, 1, \underbrace{0, \ldots, 0}_{p-k})^{T}
\end{aligned}
$$

and $X_{x}$ and $W_{x}$ are defined in (5.54).
Warning! Note that $\widehat{r}_{n}^{(k)}(x)$ is not equal to the $k^{\text {th }}$ derivative of $\widehat{r}_{n}$.
To avoid boundary and design bias, take the order of the polynomial $p$ so that $p-k$ is odd. A reasonable default is $p=k+1$. So, for estimating the first derivative, we would use local quadratic regression instead of local linear regression. The next theorem gives the large sample behavior of $\widehat{r}_{n}^{(k)}$. A proof can be found in Fan (1992) and Fan and Gijbels (1996). To state the theorem, we need a few definitions. Let $\mu_{j}=\int u^{j} K(u) d u$ and $\nu_{j}=\int u^{j} K^{2}(u) d u$. Define $(p+1) \times(p+1)$ matrices $S$ and $S^{*}$ whose $(r, s)$ entries are

$$
S_{r s}=\mu_{r+s-2}, \quad S_{r s}^{*}=\nu_{r+s-2}
$$

Also, let $c_{p}=\left(\mu_{p+1}, \ldots, \mu_{2 p+1}\right)^{T}$ and $\widetilde{c}_{p}=\left(\mu_{p+2}, \ldots, \mu_{2 p+2}\right)^{T}$. Finally, let

$$
e_{k+1}=(\underbrace{0, \ldots, 0}_{k}, 1, \underbrace{0, \ldots, 0}_{p-k})^{T} .
$$

5.144 Theorem. Let $Y_{i}=r\left(X_{i}\right)+\sigma\left(X_{i}\right) \epsilon_{i}$ for $i=1, \ldots, n$. Assume that $X_{1}, \ldots, X_{n}$ are a sample from a distribution with density $f$ and that (i) $f(x)>$ 0 , (ii) $f, r^{(p+1)}$ and $\sigma^{2}$ are continuous in a neighborhood of $x$, and (iii) $h \rightarrow 0$ and $n h \rightarrow \infty$. Then, given $X_{1}, \ldots, X_{n}$, we have the following:

$$
\begin{equation*}
\mathbb{V}\left(\widehat{r}_{n}^{(k)}(x)\right)=e_{k+1}^{T} S^{-1} S^{*} S^{-1} e_{k+1} \frac{k!^{2} \sigma^{2}(x)}{f(x) n h^{1+2 k}}+o_{P}\left(\frac{1}{n h^{1+2 k}}\right) \tag{5.145}
\end{equation*}
$$

If $p-k$ is odd, the bias is

$$
\begin{gather*}
\mathbb{E}\left(\widehat{r}_{n}^{(k)}(x)\right)-r(x)=e_{k+1}^{T} S^{-1} c_{p} \frac{k!}{(p+1)!} r^{(p+1)}(x) h^{p+1-k} \\
+o_{P}\left(h^{p+1-k}\right) \tag{5.146}
\end{gather*}
$$

If $p-k$ is even then $f^{\prime}$ and $m^{(p+2)}$ are continuous in a neighborhood ofx and $n h^{3} \rightarrow \infty$ then the bias is

$$
\begin{align*}
\mathbb{E}\left(\widehat{r}_{n}^{(k)}(x)\right)-r(x)= & e_{k+1}^{T} S^{-1} \widetilde{c}_{p} \frac{k!}{(p+2)!} \\
& \times\left(r^{(p+2)}(x)+(p+2) m^{(p+1)}(x) \frac{f^{\prime}(x)}{f(x)}\right) h^{p+2-k} \\
& +o_{P}\left(h^{p+2-k}\right) \tag{5.147}
\end{align*}
$$

Define

$$
K_{k}^{*}(t)=K(t) \sum_{\ell=0}^{p} S^{(k-1)(\ell-1)} t^{\ell}
$$

Then it can be shown that the (asymptotically) optimal bandwidth is

$$
\begin{equation*}
h_{*}=\left(\frac{C(k, p) \int \frac{\sigma^{2}(x)}{f(x)} d x}{n \int\left(r^{(p+1)}(x)\right)^{2} d x}\right)^{1 /(2 p+3)} \tag{5.148}
\end{equation*}
$$

where

$$
C(k, p)=\left(\frac{(p+1)!^{2}(2 k+1) \int K_{k}^{* 2}(t) d t}{2(p+1-k)\left(\int t^{p+1} K_{k}^{*}(t) d t\right)^{2}}\right)^{1 /(2 p+3)}
$$

Estimating a derivative is much harder then estimating the regression function because we observe the regression function (plus error) but we do not observe the derivative directly. See Chapter 6.1 of Loader (1999a) for a cogent discussion of the difficulties in estimating derivatives.

Variable Bandwidths and Adaptive Estimation. Instead of using one bandwidth $h$, we might try to use a bandwidth $h(x)$ that varies with $x$. Choosing a bandwidth this way is called variable bandwidth selection. This seems appealing because it might allow us to adapt to varying degrees of smoothness. For example, $r(x)$ might be spatially inhomogeneous meaning that it is smooth for some values of $x$ and wiggly for other values of $x$. Perhaps we should use a large bandwidth for the smooth regions and a small bandwidth for the wiggly regions. Such a procedure is said to be locally adaptive or spatially adaptive. See Chapter 4 of Fan and Gijbels (1995) and Ruppert (1997), for example. However, the improvements in the function estimate are often quite modest unless the sample size is very large and the noise level is low. For more discussion on spatial adaptation, see Chapter 9 and in particular Section 9.9.

Correlated Data. We have assumed that the errors $\epsilon_{i}=Y_{i}-r\left(x_{i}\right)$ are independent. When there is dependence between the errors, the methods need to be modified. The type of modification required depends on the type of dependence that is present. For time-ordered data, for example, time series methods are usually required. More generally, some knowledge of the dependence structure is required to devise suitable estimation methods. See Chapter 10 for more on this point.

Robustness and Quantile Regression. The estimators we have used are based on squared error loss. This is an easy loss function to use but the resulting estimator is potentially not robust to outliers. In robust regression, we choose $\widehat{a}$ to minimize

$$
\begin{equation*}
\sum_{i=1}^{n} w_{i}(x) \rho\left(\frac{\left(Y_{i}-a_{0}-a_{1}(u-x)+\cdots+\frac{a_{p}}{p!}(u-x)^{p}\right)}{s}\right) \tag{5.149}
\end{equation*}
$$

instead of minimizing (5.52). Here, $s$ is some estimate of the standard deviation of the residuals. Taking $\rho(t)=t^{2}$ gets us back to squared error loss. A more robust estimator is obtained by using Huber's function defined by the equation

$$
\begin{equation*}
\rho^{\prime}(t)=\max \{-c, \min (c, t)\} \tag{5.150}
\end{equation*}
$$

where $c$ is a tuning constant. We get back squared error as $c \rightarrow \infty$ and we get absolute error as $c \rightarrow 0$. A common choice is $c=4.685$ which provides a compromise between the two extremes. Taking

$$
\begin{equation*}
\rho(t)=|t|+(2 \alpha-1) t \tag{5.151}
\end{equation*}
$$

yields quantile regression. In this case, $\widehat{r}_{n}(x)$ estimates $\xi(x)$, where $\mathbb{P}(Y \leq$ $\xi(x) \mid X=x)=\alpha$ so that $\xi(x)$ is the $\alpha$ quantile of the conditional distribution of $Y$ given $x$. See Section 5.5 of Fan and Gijbels (1996) for more detail.

Measurement Error. In some cases we do not observe $x$ directly. Instead, we observe a corrupted version of $x$. The observed data are $\left(Y_{1}, W_{1}\right), \ldots$, $\left(Y_{n}, W_{n}\right)$ where

$$
\begin{aligned}
Y_{i} & =r\left(x_{i}\right)+\epsilon_{i} \\
W_{i} & =x_{i}+\delta_{i}
\end{aligned}
$$

for some errors $\delta_{i}$. This is called a measurement error problem or an errors-in-variables problem. Simply regressing the $Y_{i}$ 's on the $W_{i}$ 's leads to inconsistent estimates of $r(x)$. We will discuss measurement error in more detail in Chapter 10.

Dimension Reduction and Variable Selection. One way to deal with the curse of dimensionality is to try to find a low-dimension approximation
to the data. Techniques include principal component analysis, independent component analysis projection pursuit and others. See Hastie et al. (2001) for an introduction to these methods as well as relevant references.

An alternative is to perform variable selection in which covariates that do not predict $Y$ well are removed from the regression. See Zhang (1991), for example. Currently, there seems to be few variable selection methods in nonparametric regression that are both practical and have a rigorous theoretical justification.

Confidence Sets for Multiple Regression. The confidence band method in Section 5.7 can also be used for additive models. The method also extends to linear smoothers in higher dimensions as explained in Sun and Loader (1994). For more complicated methods like trees, MARS and projection pursuit regression, I am not aware of rigorous results that lead to valid bands.

Undersmoothing. One approach to dealing with the bias problem in constructing confidence sets is undersmoothing. The issue is discussed in Hall (1992b), Neumann (1995), Chen and Qin (2000), and Chen and Qin (2002). Here, we briefly discuss the results of Chen and Qin.

Let $\widehat{r}_{n}(x)$ be the local linear estimator using bandwidth $h$ and assume that the kernel $K$ has support on $[-1,1]$. Let

$$
\begin{gathered}
\alpha_{j}(x / h)=\int_{-1}^{x / h} u^{j} K(u) d u \\
\widehat{f}_{0}(x)=\frac{1}{n} \sum_{i=1}^{n} \frac{1}{h} K\left(\frac{x-X_{i}}{h}\right), \quad \widehat{f}(x)=\widehat{f}_{0}(x) / \alpha_{0}(x / h), \\
\widehat{\sigma}^{2}(x)=\frac{\frac{1}{n} \sum_{i=1}^{n} \frac{1}{h} K\left(\frac{x-X_{i}}{h}\right)\left(Y_{i}-\widehat{r}_{n}(x)\right)^{2}}{\widehat{f}_{0}(x)}
\end{gathered}
$$

and

$$
I(x)=\widehat{r}_{n}(x) \pm z_{\alpha} \sqrt{\frac{v(x / h) \widehat{\sigma}(x)}{n h \widehat{f}(x)}}
$$

where

$$
v(x / h)=\frac{\int_{-1}^{x / h}\left(\alpha_{2}(x / h)-u \alpha_{1}(x / h)\right)^{2} K^{2}(u) d u}{\left(\alpha_{0}(x / h) \alpha_{2}(x / h)-\alpha_{1}^{2}(x / h)\right)^{2}} .
$$

Undersmoothing eliminates the asymptotic bias; this means we must take $n h^{5} \rightarrow 0$. Assuming we do take $n h^{5} \rightarrow 0$, and subject to some regularity conditions, Chen and Qin (2002) show that

$$
\begin{equation*}
\mathbb{P}(r(x) \in I(x))=1-\alpha+O\left(n h^{5}+h^{2}+\frac{1}{n h}\right) \tag{5.152}
\end{equation*}
$$

at interior points and

$$
\begin{equation*}
\mathbb{P}(r(x) \in I(x))=1-\alpha+O\left(n h^{5}+h+\frac{1}{n h}\right) \tag{5.153}
\end{equation*}
$$

near the boundary. It is interesting that local linear regression eliminates boundary bias of $\widehat{r}_{n}$ but the coverage probability is poor near the boundaries. The lack of uniformity of the accuracy of the coverage probability can be fixed using an approach in Chen and Qin (2000). The confidence interval they propose is

$$
\begin{equation*}
\left\{\theta: \ell(\theta) \leq c_{\alpha}\right\} \tag{5.154}
\end{equation*}
$$

where $c_{\alpha}$ is the upper $\alpha$ quantile of a $\chi_{1}^{2}$ random variable,

$$
\ell(\theta)=2 \sum_{i=1}^{n} \log \left(1+\lambda(\theta) W_{i}\left(Y_{i}-\theta\right)\right)
$$

$\lambda(\theta)$ is defined by

$$
\begin{gathered}
\sum_{i=1}^{n} \frac{W_{i}\left(Y_{i}-\theta\right)}{1+\lambda(\theta) W_{i}\left(Y_{i}-\theta\right)}=0 \\
W_{i}=K\left(\frac{x-X_{i}}{h}\right)\left(s_{n, 2}-\frac{\left(x-X_{i}\right) s_{n, 1}}{h}\right),
\end{gathered}
$$

and

$$
s_{n, j}=\frac{1}{n h} \sum_{i=1}^{n} \frac{K\left(\frac{x-X_{i}}{h}\right)\left(x-X_{i}\right)^{j}}{h^{j}} .
$$

Assuming we do take $h n^{5} \rightarrow 0$, and subject to some regularity conditions, Chen and Qin (2000) show that

$$
\begin{equation*}
\mathbb{P}(r(x) \in I(x))=1-\alpha+O\left(n h^{5}+h^{2}+\frac{1}{n h}\right) \tag{5.155}
\end{equation*}
$$

over all $x$. The optimal bandwidth, in terms of minimizing the coverage error is

$$
h_{*}=\frac{c}{n^{1 / 3}} .
$$

Unfortunately, the constant $c$ depends on the unknown function $r$. It would appear that practical implementation is still an open research problem.

### 5.14 Bibliographic Remarks

The literature on nonparametric regression is very large. Some good starting points are Fan and Gijbels (1996), Härdle (1990), Loader (1999a), Hastie and Tibshirani (1999), and Hastie et al. (2001). A good source of the theory of splines is Wahba (1990). See also Hastie et al. (2001) and Ruppert et al. (2003). Local regression and local likelihood are discussed in detail in Loader (1999a) and Fan and Gijbels (1996). Variable bandwidth selection is discussed in Fan and Gijbels (1995).

### 5.15 Appendix

Derivation of the Tube Formula (5.100). Let $W(x)=\sum_{i=1}^{n} Z_{i} T_{i}(x)$ and recall that $\|T(x)\|^{2}=\sum_{i=1}^{n} T_{i}(x)^{2}=1$ so that the vector $T(x)$ is on the unit sphere, for each $x$. Since, $Z=\left(Z_{1}, \ldots, Z_{n}\right)$ are multivariate Normal,

$$
\begin{aligned}
\mathbb{P}\left(\sup _{x} W(x)>c\right) & =\mathbb{P}\left(\sup _{x}\langle Z, T(x)\rangle>c\right) \\
& =\mathbb{P}\left(\sup _{x}\left\langle\frac{Z}{\|Z\|}, T(x)\right\rangle>\frac{c}{\|Z\|}\right) \\
& =\int_{c^{2}}^{\infty} \mathbb{P}\left(\sup _{x}\langle U, T(x)\rangle>\frac{c}{\sqrt{y}}\right) h(y) d y
\end{aligned}
$$

where $U=\left(U_{1}, \ldots, U_{n}\right)$ is uniformly distributed on the $n$-1-dimensional unit sphere $S$ and $h(y)$ is the density for a $\chi^{2}$ with $n$ degrees of freedom. Since $\|U-T(x)\|^{2}=2(1-\langle U, T(x)\rangle)$, we see that $\sup _{x}\langle U, T(x)\rangle>\frac{c}{\sqrt{y}}$ if and only if $U \in \operatorname{tube}(r, M)$ where $r=\sqrt{2(1-c / \sqrt{y})}, M=\{T(x): x \in \mathcal{X}\}$ is a manifold on the sphere $S$,

$$
\operatorname{tube}(r, M)=\{u: d(u, M) \leq r\}
$$

and

$$
d(u, M)=\inf _{x \in \mathcal{X}}\|u-T(x)\|
$$

Therefore,

$$
\begin{aligned}
\mathbb{P}\left(\sup _{x}\langle U, T(x)\rangle>\frac{c}{\sqrt{y}}\right) & =\mathbb{P}(U \in \operatorname{tube}(r, M)) \\
& =\frac{\text { volume }(\operatorname{tube}(r, M))}{A_{n}}
\end{aligned}
$$

where $A_{n}=2 \pi^{n / 2} / \Gamma(n / 2)$ is the area of the unit sphere. The formula for volume(tube ( $r, M$ )) was derived by Hotelling (1939) and Naiman (1990) and
is given by

$$
\kappa_{0} \frac{A_{n}}{A_{2}} \mathbb{P}\left(B_{1,(n-2) / 2} \geq w^{2}\right)+\ell_{0} \frac{A_{n}}{2 A_{1}} \mathbb{P}\left(B_{1 / 2,(n-1) / 2} \geq w^{2}\right)
$$

where $w=c / \sqrt{y}$. Inserting this into the integral and ignoring terms of order smaller than $c^{-1 / 2} e^{-c^{2} / 2}$ yields (5.100).

The formula can also be obtained using the upcrossing theory of Rice (1939). Specifically, if $W$ is a Gaussian process on $[0,1]$ and if $N_{c}$ denotes the number of upcrossings of $W$ above $c$ then

$$
\begin{aligned}
\mathbb{P}\left(\sup _{x} W(x)>c\right) & =\mathbb{P}\left(N_{c} \geq 1 \text { or } W(0)>c\right) \\
& \leq \mathbb{P}\left(N_{c} \geq 1\right)+\mathbb{P}(W(0)>c) \\
& \leq \mathbb{E}\left(N_{c}\right)+\mathbb{P}(W(0)>c)
\end{aligned}
$$

Since $W(0)$ has a Normal distribution, the second term can be easily computed. Moreover, under smoothness conditions on $W$ we have

$$
\begin{equation*}
\mathbb{E}\left(N_{c}\right)=\int_{0}^{1} \int_{0}^{\infty} y p_{t}(c, y) d y d t \tag{5.156}
\end{equation*}
$$

where $p_{t}$ is the density of $\left(W(t), W^{\prime}(t)\right)$.

### 5.16 Exercises

1. In Example 5.24, construct the smoothing matrix $L$ and verify that $\nu=m$.
2. Prove Theorem 5.34.
3. Get the data on fragments of glass collected in forensic work from the book website. Let $Y$ be refractive index and let $x$ be aluminium content (the fourth variable). Perform a nonparametric regression to fit the model $Y=r(x)+\epsilon$. Use the following estimators: (i) regressogram, (ii) kernel, (iii) local linear, (iv) spline. In each case, use cross-validation to choose the amount of smoothing. Estimate the variance. Construct 95 percent confidence bands for your estimates. Pick a few values of $x$ and, for each value, plot the effective kernel for each smoothing method. Visually compare the effective kernels.
4. Get the motorcycle data from the book website. The covariate is time (in milliseconds) and the response is acceleration at time of impact. Use cross-validation to fit a smooth curve using local linear regression.
5. Show that with suitable smoothness assumptions on $r(x), \widehat{\sigma}^{2}$ in equation (5.89) is a consistent estimator of $\sigma^{2}$.
6. Prove Theorem 5.34.
7. Prove Theorem 5.60.
8. Find conditions under which the estimate in equation (5.86) is consistent.
9. Consider the data in Exercise 3. Examine the fit as a function of the bandwidth $h$. Do this by plotting the fit for many values of $h$. Add confidence bands to all the fits. If you are feeling very ambitious, read Chaudhuri and Marron (1999) and apply that method.
10. Using five equally spaced knots on $(0,1)$, construct a B-spline basis of order $M$ for $M=1, \ldots, 5$. Plot the basis functions.
11. Get the motorcycle data from the book website. Fit a cubic regression spline with equally spaced knots. Use leave-one-out cross-validation to choose the number of knots. Now fit a smoothing spline and compare the fits.
12. Recall the Doppler function defined in Example 5.63. Generate 1000 observations from the model $Y_{i}=r\left(x_{i}\right)+\sigma \epsilon_{i}$ where $x_{i}=i / n$ and $\epsilon_{i} \sim N(0,1)$. Make three data sets corresponding to $\sigma=.1, \sigma=1$ and $\sigma=3$. Plot the data. Estimate the function using local linear regression. Plot the cross-validation score versus the bandwidth. Plot the fitted function. Find and plot a 95 percent confidence band.
13. Repeat the previous question but use smoothing splines.
14. Download the air quality data set from the book website. Model ozone as a function of temperature. Use kernel regression and compare the fit you get when you choose the bandwidth using cross-validation, generalized cross-validation, $C_{p}$ and the plug-in method.
15. Let $Y_{i} \sim N\left(\mu_{i}, 1\right)$ for $i=1, \ldots, n$ be independent observations. Find the estimators that minimizes each of the following penalized sums of squares:
(a) $\quad \sum_{i=1}^{n}\left(Y_{i}-\widehat{\mu}_{i}\right)^{2}+\lambda \sum_{i=1}^{n} \widehat{\mu}_{i}^{2}$

$$
\begin{array}{ll}
\text { (b) } & \sum_{i=1}^{n}\left(Y_{i}-\widehat{\mu}_{i}\right)^{2}+\lambda \sum_{i=1}^{n}\left|\widehat{\mu}_{i}\right| \\
\text { (c) } & \sum_{i=1}^{n}\left(Y_{i}-\widehat{\mu}_{i}\right)^{2}+\lambda \sum_{i=1}^{n} I\left(\widehat{\mu}_{i}=0\right) .
\end{array}
$$

16. Show that a locally polynomial smoother of order $p$ reproduces polynomials of order $p$.
17. Suppose that $r:[0,1] \rightarrow \mathbb{R}$ satisfies the following Lipschitz condition:

$$
\begin{equation*}
\sup _{0 \leq x<y \leq 1}|r(y)-r(x)| \leq L|y-x| \tag{5.157}
\end{equation*}
$$

where $L>0$ is given. The class of all such functions is denoted by $\mathcal{F}_{\text {lip }}(L)$. What is the maximum bias of a kernel estimator $\widehat{r}_{n}$ based on bandwidth $h$, if $r \in \mathcal{F}_{\text {lip }}(L)$ ?
18. Implement quantile regression on the glass data (Exercise 3) with $\alpha=1 / 2$.
19. Prove that the weights $\ell_{i}(x)$ for the local polynomial smoother satisfy

$$
\begin{equation*}
\ell_{i}(x)=K\left(\frac{x_{i}-x}{h}\right) P_{i}(x) \tag{5.158}
\end{equation*}
$$

for some polynomial

$$
P_{i}(x)=\alpha_{0}+\alpha_{1}\left(x_{i}-x\right)+\cdots+\alpha_{p}\left(x_{i}-x\right)^{p} .
$$

Moreover, if the $i^{\text {th }}$ observation $\left(x_{i}, Y_{i}\right)$ is omitted, the resulting weights satisfy (5.32). Thus, while we took (5.32) as the definition of the leave-one-out weights, one can derive this form of the weights.
20. Suppose that $\ell_{i}(x)=K\left(\left(x-x_{i}\right) / h\right)$ for some smooth kernel $K$ and that the $x_{i}$ 's are equally spaced. Define $\kappa_{0}$ as in (5.101). Show that, if we ignore boundary effects,

$$
\kappa_{0} \approx\left(\frac{b-a}{h}\right) \frac{\left\|K^{\prime}\right\|}{\|K\|}
$$

where $\|g\|^{2}=\int_{a}^{b} g^{2}(x) d x$.
21. Show how to construct a confidence band for the derivative estimator $\widehat{r}^{(k)}$ given in (5.143). Hint: Note that the estimator is linear and follow the construction of the confidence band for $\widehat{r}_{n}(x)$.
22. Download the air quality data set from the book website. Model ozone as a function of solar R , wind and temperature. Use (i) multiple local linear regression, (ii) projection pursuit, (iii) additive regression, (iv) regression trees, and (v) MARS. Compare the results.
23. Explain how to construct confidence bands in additive models. Apply this to the data from Exercise 22.
24. Let $\widehat{r}_{n}\left(x_{1}, x_{2}\right)=\sum_{i=1}^{n} Y_{i} \ell_{i}\left(x_{1}, x_{2}\right)$ be a linear estimator of the multiple regression function $r\left(x_{1}, x_{2}\right)$. Suppose we want to test the hypothesis that the covariate $x_{2}$ can be dropped from the regression. One possibility is to form a linear estimator of the form $\widetilde{r}_{n}\left(x_{1}\right)=\sum_{i=1}^{n} Y_{i} \widetilde{\ell}_{i}\left(x_{1}\right)$ and then compute

$$
T=\sum_{i=1}^{n}\left(\widehat{r}_{n}\left(x_{1 i}, x_{2 i}\right)-\widetilde{r}_{n}\left(x_{1 i}\right)\right)^{2}
$$

(i) Assume that the true model is $Y_{i}=r\left(x_{1 i}\right)+\epsilon_{i}$ where $\epsilon_{i} \sim N\left(0, \sigma^{2}\right)$. For simplicity take $\sigma$ known. Find an expression for the distribution of $T$.
(ii) The null distribution in part (i) depends on the unknown function $r\left(x_{1}\right)$. How might you estimate the null distribution?
(iii) Create simulated data from the model in (i) (use any function $r\left(x_{1}\right)$ you like) and see if your proposed method in (ii) does approximate the null distribution.

## 6

## Density Estimation

Let $F$ be a distribution with probability density $f=F^{\prime}$ and let

$$
X_{1}, \ldots, X_{n} \sim F
$$

be an IID sample from $F$. The goal of nonparametric density estimation is to estimate $f$ with as few assumptions about $f$ as possible. We denote the estimator by $\widehat{f}_{n}$. As with nonparametric regression, the estimator will depend on a smoothing parameter $h$ and choosing $h$ carefully is important.
6.1 Example (Bart Simpson). The top left plot in Figure 6.1 shows the density

$$
\begin{equation*}
f(x)=\frac{1}{2} \phi(x ; 0,1)+\frac{1}{10} \sum_{j=0}^{4} \phi(x ;(j / 2)-1,1 / 10) \tag{6.2}
\end{equation*}
$$

where $\phi(x ; \mu, \sigma)$ denotes a Normal density with mean $\mu$ and standard deviation $\sigma$. Marron and Wand (1992) call this density "the claw" although we will call it the Bart Simpson density. Based on 1000 draws from $f$, I computed a kernel density estimator, described later in the chapter. The top right plot is based on a small bandwidth $h$ which leads to undersmoothing. The bottom right plot is based on a large bandwidth $h$ which leads to oversmoothing. The bottom left plot is based on a bandwidth $h$ which was chosen to minimize estimated risk. This leads to a much more reasonable density estimate.


FIGURE 6.1. The Bart Simpson density from Example 6.1. Top left: true density. The other plots are kernel estimators based on $n=1000$ draws. Bottom left: bandwidth $h=0.05$ chosen by leave-one-out cross-validation. Top right: bandwidth $h / 10$. Bottom right: bandwidth 10 h .

### 6.1 Cross-Validation

We will evaluate the quality of an estimator $\widehat{f}_{n}$ with the risk, or integrated mean squared error, $R=\mathbb{E}(L)$ where

$$
L=\int\left(\widehat{f}_{n}(x)-f(x)\right)^{2} d x
$$

is the integrated squared error loss function. The estimators will depend on some smoothing parameter $h$ and we will choose $h$ to minimize an estimate of the risk. The usual method for estimating risk is leave-one-out cross-validation. The details are different for density estimation than for regression. In the regression case, the cross-validation score was defined as $\sum_{i=1}^{n}\left(Y_{i}-\widehat{r}_{(-i)}\left(x_{i}\right)\right)^{2}$ but in density estimation, there is no response variable $Y$. Instead, we proceed as follows.

The loss function, which we now write as a function of $h$, (since $\widehat{f}_{n}$ will depend on some smoothing parameter $h$ ) is

$$
\begin{aligned}
L(h) & =\int\left(\widehat{f}_{n}(x)-f(x)\right)^{2} d x \\
& =\int \widehat{f}_{n}^{2}(x) d x-2 \int \widehat{f}_{n}(x) f(x) d x+\int f^{2}(x) d x
\end{aligned}
$$

The last term does not depend on $h$ so minimizing the loss is equivalent to minimizing the expected value of

$$
\begin{equation*}
J(h)=\int \widehat{f}_{n}^{2}(x) d x-2 \int \widehat{f}_{n}(x) f(x) d x \tag{6.3}
\end{equation*}
$$

We shall refer to $\mathbb{E}(J(h))$ as the risk, although it differs from the true risk by the constant term $\int f^{2}(x) d x$.

### 6.4 Definition. The cross-validation estimator of risk is

$$
\begin{equation*}
\widehat{J}(h)=\int\left(\widehat{f}_{n}(x)\right)^{2} d x-\frac{2}{n} \sum_{i=1}^{n} \widehat{f}_{(-i)}\left(X_{i}\right) \tag{6.5}
\end{equation*}
$$

where $\widehat{f}_{(-i)}$ is the density estimator obtained after removing the $i^{\text {th }}$ observation. We refer to $\widehat{J}(h)$ as the cross-validation score or estimated risk.

### 6.2 Histograms

Perhaps the simplest nonparametric density estimator is the histogram. Suppose $f$ has its support on some interval which, without loss of generality, we take to be $[0,1]$. Let $m$ be an integer and define bins

$$
\begin{equation*}
B_{1}=\left[0, \frac{1}{m}\right), \quad B_{2}=\left[\frac{1}{m}, \frac{2}{m}\right), \quad \ldots, \quad B_{m}=\left[\frac{m-1}{m}, 1\right] . \tag{6.6}
\end{equation*}
$$

Define the binwidth $h=1 / m$, let $Y_{j}$ be the number of observations in $B_{j}$, let $\widehat{p}_{j}=Y_{j} / n$ and let $p_{j}=\int_{B_{j}} f(u) d u$.

The histogram estimator is defined by

$$
\begin{equation*}
\widehat{f}_{n}(x)=\sum_{j=1}^{m} \frac{\widehat{p}_{j}}{h} I\left(x \in B_{j}\right) . \tag{6.7}
\end{equation*}
$$

To understand the motivation for this estimator, note that, for $x \in B_{j}$ and $h$ small,

$$
\mathbb{E}\left(\widehat{f}_{n}(x)\right)=\frac{\mathbb{E}\left(\widehat{p}_{j}\right)}{h}=\frac{p_{j}}{h}=\frac{\int_{B_{j}} f(u) d u}{h} \approx \frac{f(x) h}{h}=f(x) .
$$

6.8 Example. Figure 6.2 shows three different histograms based on $n=1,266$ data points from an astronomical sky survey. These are the data from Example 4.3. Each data point represents a "redshift," roughly speaking, the distance from us to a galaxy. Choosing the right number of bins involves finding a good tradeoff between bias and variance. We shall see later that the top left histogram has too many bins resulting in oversmoothing and too much bias. The bottom left histogram has too few bins resulting in undersmoothing. The top right histogram is based on 308 bins (chosen by cross-validation). The histogram reveals the presence of clusters of galaxies.


FIGURE 6.2. Three versions of a histogram for the astronomy data. The top left histogram has too many bins. The bottom left histogram has too few bins. The top right histogram uses 308 bins (chosen by cross-validation). The lower right plot shows the estimated risk versus the number of bins.
6.9 Theorem. Consider fixed $x$ and fixed $m$, and let $B_{j}$ be the bin containing $x$. Then,

$$
\begin{equation*}
\mathbb{E}\left(\widehat{f}_{n}(x)\right)=\frac{p_{j}}{h} \quad \text { and } \quad \mathbb{V}\left(\widehat{f}_{n}(x)\right)=\frac{p_{j}\left(1-p_{j}\right)}{n h^{2}} \tag{6.10}
\end{equation*}
$$

6.11 Theorem. Suppose that $f^{\prime}$ is absolutely continuous and that $\int\left(f^{\prime}(u)\right)^{2} d u<\infty$. Then

$$
\begin{equation*}
R\left(\widehat{f}_{n}, f\right)=\frac{h^{2}}{12} \int\left(f^{\prime}(u)\right)^{2} d u+\frac{1}{n h}+o\left(h^{2}\right)+o\left(\frac{1}{n}\right) . \tag{6.12}
\end{equation*}
$$

The value $h^{*}$ that minimizes (6.12) is

$$
\begin{equation*}
h^{*}=\frac{1}{n^{1 / 3}}\left(\frac{6}{\int\left(f^{\prime}(u)\right)^{2} d u}\right)^{1 / 3} \tag{6.13}
\end{equation*}
$$

With this choice of binwidth,

$$
\begin{equation*}
R\left(\widehat{f}_{n}, f\right) \sim \frac{C}{n^{2 / 3}} \tag{6.14}
\end{equation*}
$$

where $C=(3 / 4)^{2 / 3}\left(\int\left(f^{\prime}(u)\right)^{2} d u\right)^{1 / 3}$.
The proof of Theorem 6.11 is in the appendix. We see that with an optimally chosen binwidth, the risk decreases to 0 at rate $n^{-2 / 3}$. We will see shortly that kernel estimators converge at the faster rate $n^{-4 / 5}$ and that, in a certain sense, no faster rate is possible; see Theorem 6.31. The formula for the optimal binwidth $h^{*}$ is of theoretical interest but it is not useful in practice since it depends on the unknown function $f$. In practice, we use cross-validation as described in Section 6.1. There is a simple formula for computing the crossvalidation score $\widehat{J}(h)$.
6.15 Theorem. The following identity holds:

$$
\begin{equation*}
\widehat{J}(h)=\frac{2}{h(n-1)}-\frac{n+1}{h(n-1)} \sum_{j=1}^{m} \widehat{p}_{j}^{2} . \tag{6.16}
\end{equation*}
$$

6.17 Example. We used cross-validation in the astronomy example. We find that $m=308$ is an approximate minimizer. The histogram in the top right plot in Figure 6.2 was constructed using $m=308$ bins. The bottom right plot shows the estimated risk, or more precisely, $\widehat{J}$, plotted versus the number of bins.

Next, we want a confidence set for $f$. Suppose $\widehat{f}_{n}$ is a histogram with $m$ bins and binwidth $h=1 / m$. For reasons explained in Section 5.7, it is difficult to construct a confidence set for $f$. Instead, we shall make confidence statements about $f$ at the resolution of the histogram. Thus, define

$$
\begin{equation*}
\bar{f}_{n}(x)=\mathbb{E}\left(\widehat{f}_{n}(x)\right)=\sum_{j=1}^{m} \frac{p_{j}}{h} I\left(x \in B_{j}\right) \tag{6.18}
\end{equation*}
$$

where $p_{j}=\int_{B_{j}} f(u) d u$. Think of $\bar{f}_{n}(x)$ as a "histogramized" version of $f$. Recall that a pair of functions $(\ell, u)$ is a $1-\alpha$ confidence band for $\bar{f}_{n}$ if

$$
\begin{equation*}
\mathbb{P}\left(\ell(x) \leq \bar{f}_{n}(x) \leq u(x) \quad \text { for all } x\right) \geq 1-\alpha \tag{6.19}
\end{equation*}
$$

We could use the type of reasoning as in (5.100) but, instead, we take a simpler route.
6.20 Theorem. Let $m=m(n)$ be the number of bins in the histogram $\widehat{f}_{n}$. Assume that $m(n) \rightarrow \infty$ and $m(n) \log n / n \rightarrow 0$ as $n \rightarrow \infty$. Define

$$
\begin{align*}
& \ell_{n}(x)=\left(\max \left\{\sqrt{\widehat{f}_{n}(x)}-c, 0\right\}\right)^{2} \\
& u_{n}(x)=\left(\sqrt{\widehat{f}_{n}(x)}+c\right)^{2} \tag{6.21}
\end{align*}
$$

where

$$
\begin{equation*}
c=\frac{z_{\alpha /(2 m)}}{2} \sqrt{\frac{m}{n}} . \tag{6.22}
\end{equation*}
$$

Then, $\left(\ell_{n}(x), u_{n}(x)\right)$ is an approximate $1-\alpha$ confidence band for $\bar{f}_{n}$.
Proof. Here is an outline of the proof. From the central limit theorem, and assuming $1-p_{j} \approx 1, \widehat{p}_{j} \approx N\left(p_{j}, p_{j}\left(1-p_{j}\right) / n\right)$. By the delta method, $\sqrt{\widehat{p}_{j}} \approx N\left(\sqrt{p_{j}}, 1 /(4 n)\right)$. Moreover, the $\sqrt{\widehat{p}_{j}}$ 's are approximately independent. Therefore,

$$
\begin{equation*}
2 \sqrt{n}\left(\sqrt{\widehat{p}_{j}}-\sqrt{p}_{j}\right) \approx Z_{j} \tag{6.23}
\end{equation*}
$$

where $Z_{1}, \ldots, Z_{m} \sim N(0,1)$. Let

$$
A=\left\{\ell_{n}(x) \leq \bar{f}_{n}(x) \leq u_{n}(x) \text { for all } x\right\}=\left\{\max _{x}\left|\sqrt{\widehat{f}_{n}(x)}-\sqrt{\bar{f}(x)}\right| \leq c\right\}
$$

Then,

$$
\mathbb{P}\left(A^{c}\right)=\mathbb{P}\left(\max _{x}\left|\sqrt{\widehat{f}_{n}(x)}-\sqrt{\bar{f}(x)}\right|>c\right)
$$



FIGURE 6.3. Ninety-five percent confidence envelope for astronomy data using $m=308$ bins.

$$
\begin{aligned}
& =\mathbb{P}\left(\max _{j} 2 \sqrt{n}\left|\sqrt{\widehat{p}_{j}}-\sqrt{p_{j}}\right|>z_{\alpha /(2 m)}\right) \\
& \approx \mathbb{P}\left(\max _{j}\left|Z_{j}\right|>z_{\alpha /(2 m)}\right) \leq \sum_{j=1}^{m} \mathbb{P}\left(\left|Z_{j}\right|>z_{\alpha /(2 m)}\right) \\
& =\sum_{j=1}^{m} \frac{\alpha}{m}=\alpha .
\end{aligned}
$$

6.24 Example. Figure 6.3 shows a 95 percent confidence envelope for the astronomy data. We see that even with over 1000 data points, there is still substantial uncertainty about $f$ as reflected by the wide bands.

### 6.3 Kernel Density Estimation

Histograms are not smooth. In this section we discuss kernel density estimators which are smoother and which converge to the true density faster. Recall that the word kernel refers to any smooth function $K$ satisfying the conditions given in (4.22). See Section 4.2 for examples of kernels.
6.25 Definition. Given a kernel $K$ and a positive number $h$, called the bandwidth, the kernel density estimator is defined to be

$$
\begin{equation*}
\widehat{f}_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} \frac{1}{h} K\left(\frac{x-X_{i}}{h}\right) \tag{6.26}
\end{equation*}
$$

This amounts to placing a smoothed out lump of mass of size $1 / n$ over each data point $X_{i}$; see Figure 6.4.


FIGURE 6.4. A kernel density estimator $\widehat{f}_{n}$. At each point $x, \widehat{f}_{n}(x)$ is the average of the kernels centered over the data points $X_{i}$. The data points are indicated by short vertical bars. The kernels are not drawn to scale.

As with kernel regression, the choice of kernel $K$ is not crucial, but the choice of bandwidth $h$ is important. Figure 6.5 shows density estimates with several different bandwidths. (This is the same as Figure 4.3.) Look also at Figure 6.1. We see how sensitive the estimate $\widehat{f}_{n}$ is to the choice of $h$. Small bandwidths give very rough estimates while larger bandwidths give smoother estimates. In general we will let the bandwidth depend on the sample size so we write $h_{n}$. Here are some properties of $\widehat{f}_{n}$.


FIGURE 6.5. Kernel density estimators and estimated risk for the astronomy data. Top left: oversmoothed. Top right: just right (bandwidth chosen by cross-validation). Bottom left: undersmoothed. Bottom right: cross-validation curve as a function of bandwidth $h$. The bandwidth was chosen to be the value of $h$ where the curve is a minimum.
6.27 Theorem. Assume that $f$ is continuous at $x$ and that $h_{n} \rightarrow 0$ and $n h_{n} \rightarrow$ $\infty$ as $n \rightarrow \infty$. Then $\widehat{f}_{n}(x) \xrightarrow{\mathrm{P}} f(x)$.
6.28 Theorem. Let $R_{x}=\mathbb{E}(f(x)-\widehat{f}(x))^{2}$ be the risk at a point $x$ and let $R=$ $\int R_{x} d x$ denote the integrated risk. Assume that $f^{\prime \prime}$ is absolutely continuous and that $\int\left(f^{\prime \prime \prime}(x)\right)^{2} d x<\infty$. Also, assume that $K$ satisfies (4.22). Then,

$$
R_{x}=\frac{1}{4} \sigma_{K}^{4} h_{n}^{4}\left(f^{\prime \prime}(x)\right)^{2}+\frac{f(x) \int K^{2}(x) d x}{n h_{n}}+O\left(\frac{1}{n}\right)+O\left(h_{n}^{6}\right)
$$

and

$$
\begin{equation*}
R=\frac{1}{4} \sigma_{K}^{4} h_{n}^{4} \int\left(f^{\prime \prime}(x)\right)^{2} d x+\frac{\int K^{2}(x) d x}{n h}+O\left(\frac{1}{n}\right)+O\left(h_{n}^{6}\right) \tag{6.29}
\end{equation*}
$$

where $\sigma_{K}^{2}=\int x^{2} K(x) d x$.

Proof. Write $K_{h}(x, X)=h^{-1} K((x-X) / h)$ and $\widehat{f}_{n}(x)=n^{-1} \sum_{i} K_{h}\left(x, X_{i}\right)$. Thus, $\mathbb{E}\left[\widehat{f}_{n}(x)\right]=\mathbb{E}\left[K_{h}(x, X)\right]$ and $\mathbb{V}\left[\widehat{f}_{n}(x)\right]=n^{-1} \mathbb{V}\left[K_{h}(x, X)\right]$. Now,

$$
\begin{aligned}
\mathbb{E}\left[K_{h}(x, X)\right] & =\int \frac{1}{h} K\left(\frac{x-t}{h}\right) f(t) d t \\
& =\int K(u) f(x-h u) d u \\
& =\int K(u)\left[f(x)-h u f^{\prime}(x)+\frac{h^{2} u^{2}}{2} f^{\prime \prime}(x)+\cdots\right] d u \\
& =f(x)+\frac{1}{2} h^{2} f^{\prime \prime}(x) \int u^{2} K(u) d u \cdots
\end{aligned}
$$

since $\int K(x) d x=1$ and $\int x K(x) d x=0$. The bias is

$$
\mathbb{E}\left[K_{h_{n}}(x, X)\right]-f(x)=\frac{1}{2} \sigma_{K}^{2} h_{n}^{2} f^{\prime \prime}(x)+O\left(h_{n}^{4}\right)
$$

By a similar calculation,

$$
\mathbb{V}\left[\widehat{f}_{n}(x)\right]=\frac{f(x) \int K^{2}(x) d x}{n h_{n}}+O\left(\frac{1}{n}\right)
$$

The first result then follows since the risk is the squared bias plus variance. The second result follows from integrating the first.

If we differentiate (6.29) with respect to $h$ and set it equal to 0 , we see that the asymptotically optimal bandwidth is

$$
\begin{equation*}
h_{*}=\left(\frac{c_{2}}{c_{1}^{2} A(f) n}\right)^{1 / 5} \tag{6.30}
\end{equation*}
$$

where $c_{1}=\int x^{2} K(x) d x, c_{2}=\int K(x)^{2} d x$ and $A(f)=\int\left(f^{\prime \prime}(x)\right)^{2} d x$. This is informative because it tells us that the best bandwidth decreases at rate $n^{-1 / 5}$. Plugging $h_{*}$ into (6.29), we see that if the optimal bandwidth is used then $R=O\left(n^{-4 / 5}\right)$. As we saw, histograms converge at rate $O\left(n^{-2 / 3}\right)$ showing that kernel estimators are superior in rate to histograms. According to the next theorem, there does not exist an estimator that converges faster than $O\left(n^{-4 / 5}\right)$. For a proof, see, for example, Chapter 24 of van der Vaart (1998).
6.31 Theorem. Let $\mathcal{F}$ be the set of all probability density functions and let $f^{(m)}$ denote the $m^{\text {th }}$ derivative of $f$. Define

$$
\mathcal{F}_{m}(c)=\left\{f \in \mathcal{F}: \quad \int\left|f^{(m)}(x)\right|^{2} d x \leq c^{2}\right\}
$$

For any estimator $\widehat{f}_{n}$,

$$
\begin{equation*}
\sup _{f \in \mathcal{F}_{m}(c)} \mathbb{E}_{f} \int\left(\widehat{f}_{n}(x)-f(x)\right)^{2} d x \geq b\left(\frac{1}{n}\right)^{2 m /(2 m+1)} \tag{6.32}
\end{equation*}
$$

where $b>0$ is a universal constant that depends only on $m$ and $c$.
In particular, taking $m=2$ in the previous theorem we see that $n^{-4 / 5}$ is the fastest possible rate.

In practice, the bandwidth can be chosen by cross-validation but first we describe another method which is sometimes used when $f$ is thought to be very smooth. Specifically, we compute $h_{*}$ from (6.30) under the idealized assumption that $f$ is Normal. This yields $h_{*}=1.06 \sigma n^{-1 / 5}$. Usually, $\sigma$ is estimated by $\min \{s, Q / 1.34\}$ where $s$ is the sample standard deviation and $Q$ is the interquartile range. ${ }^{1}$ This choice of $h_{*}$ works well if the true density is very smooth and is called the Normal reference rule.

## The Normal Reference Rule

For smooth densities and a Normal kernel, use the bandwidth

$$
h_{n}=\frac{1.06 \widehat{\sigma}}{n^{1 / 5}}
$$

where

$$
\widehat{\sigma}=\min \left\{s, \frac{Q}{1.34}\right\}
$$

Since we don't want to necessarily assume that $f$ is very smooth, it is usually better to estimate $h$ using cross-validation. Recall from Section 6.1 that the cross-validation score is

$$
\begin{equation*}
\widehat{J}(h)=\int \widehat{f}^{2}(x) d x-\frac{2}{n} \sum_{i=1}^{n} \widehat{f}_{-i}\left(X_{i}\right) \tag{6.33}
\end{equation*}
$$

where $\widehat{f}_{-i}$ denotes the kernel estimator obtained by omitting $X_{i}$. The next theorem gives a simpler expression for $\widehat{J}$.
6.34 Theorem. For any $h>0$,

$$
\mathbb{E}[\widehat{J}(h)]=\mathbb{E}[J(h)]
$$

[^11]Also,

$$
\begin{equation*}
\widehat{J}(h)=\frac{1}{h n^{2}} \sum_{i} \sum_{j} K^{*}\left(\frac{X_{i}-X_{j}}{h}\right)+\frac{2}{n h} K(0)+O\left(\frac{1}{n^{2}}\right) \tag{6.35}
\end{equation*}
$$

where $K^{*}(x)=K^{(2)}(x)-2 K(x)$ and $K^{(2)}(z)=\int K(z-y) K(y) d y$.
6.36 Remark. When $K$ is a $\mathrm{N}(0,1)$ Gaussian kernel then $K^{(2)}(z)$ is the $N(0,2)$ density. Also, we should point out that the estimator $\widehat{f}_{n}$ and the cross-validation score (6.35) can be computed quickly using the fast Fourier transform; see pages 61-66 of Silverman (1986).

A justification for cross-validation is given by the following remarkable theorem due to Stone (1984).
6.37 Theorem (Stone's theorem). Suppose that $f$ is bounded. Let $\widehat{f}_{h}$ denote the kernel estimator with bandwidth $h$ and let $\widehat{h}$ denote the bandwidth chosen by cross-validation. Then,

$$
\begin{equation*}
\frac{\int\left(f(x)-\widehat{f}_{\widehat{h}}(x)\right)^{2} d x}{\inf _{h} \int\left(f(x)-\widehat{f}_{h}(x)\right)^{2} d x} \stackrel{\text { a.s. }}{ } 1 \tag{6.38}
\end{equation*}
$$

The bandwidth for the density estimator in the upper right panel of Figure 6.5 is based on cross-validation. In this case it worked well but of course there are lots of examples where there are problems. Do not assume that, if the estimator $\widehat{f}$ is wiggly, then cross-validation has let you down. The eye is not a good judge of risk.

Another approach to bandwidth selection called plug-in bandwidths. The idea is as follows. The (asymptotically) optimal bandwidth is given in equation (6.30). The only unknown quantity in that formula is $A(f)=\int\left(f^{\prime \prime}(x)\right)^{2} d x$. If we have an estimate $\widehat{f}^{\prime \prime}$ of $f^{\prime \prime}$, then we can plug this estimate into the formula for the optimal bandwidth $h_{*}$. There is a rich and interesting literature on this and similar approaches. The problem with this approach is that estimating $f^{\prime \prime}$ is harder than estimating $f$. Indeed, we need to make stronger assumptions about $f$ to estimate $f^{\prime \prime}$. But if we make these stronger assumptions then the (usual) kernel estimator for $f$ is not appropriate. Loader (1999b) has investigated this issue in detail and provides evidence that the plug-in bandwidth approach might not be reliable. There are also methods that apply corrections to plug-in rules; see Hjort (1999).

A generalization of the kernel method is to use adaptive kernels where one uses a different bandwidth $h(x)$ for each point $x$. One can also use a
different bandwidth $h\left(x_{i}\right)$ for each data point. This makes the estimator more flexible and allows it to adapt to regions of varying smoothness. But now we have the very difficult task of choosing many bandwidths instead of just one. See Chapter 9 for more on adaptive methods.

Constructing confidence bands for kernel density estimators is more complicated than for regression. We discuss one possible approach in Section 6.6.

### 6.4 Local Polynomials

In Chapter 5 we saw that kernel regression suffers from boundary bias and that this bias can be reduced by using local polynomials. The same is true for kernel density estimation. But what density estimation method corresponds to local polynomial regression? One possibility, developed by Loader (1999a) and Hjort and Jones (1996) is to use local likelihood density estimation.

The usual definition of $\log$-likelihood is $\mathcal{L}(f)=\sum_{i=1}^{n} \log f\left(X_{i}\right)$. It is convenient to generalize this definition to

$$
\mathcal{L}(f)=\sum_{i=1}^{n} \log f\left(X_{i}\right)-n\left(\int f(u) d u-1\right)
$$

The second term is zero when $f$ integrates to one. Including this term allows us to maximize over all non-negative $f$ with imposing the constraint that $\int f(u) d u=1$. The local version of this log-likelihood is as follows:
6.39 Definition. Given a kernel $K$ and bandwidth $h$, the local log-likelihood at a target value $x$ is

$$
\begin{equation*}
\mathcal{L}_{x}(f)=\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) \log f\left(X_{i}\right)-n \int K\left(\frac{u-x}{h}\right) f(u) d u \tag{6.40}
\end{equation*}
$$

The above definition is for an arbitrary density $f$. We are interested in the case where we approximate $\log f(u)$ by a polynomial in a neighborhood of $x$. Thus we write

$$
\begin{equation*}
\log f(u) \approx P_{x}(a, u) \tag{6.41}
\end{equation*}
$$

where

$$
\begin{equation*}
P_{x}(a, u)=a_{0}+a_{1}(x-u)+\cdots+a_{p} \frac{(x-u)^{p}}{p!} . \tag{6.42}
\end{equation*}
$$

Plugging (6.41) into (6.40) yields the local polynomial log-likelihood

$$
\begin{equation*}
\mathcal{L}_{x}(a)=\sum_{i=1}^{n} K\left(\frac{X_{i}-x}{h}\right) P_{x}\left(a, X_{i}\right)-n \int K\left(\frac{u-x}{h}\right) e^{P_{x}(a, u)} d u . \tag{6.43}
\end{equation*}
$$

6.44 Definition. Let $\widehat{a}=\left(\widehat{a}_{0}, \ldots, \widehat{a}_{p}\right)^{T}$ maximize $\mathcal{L}_{x}(a)$. The local likelihood density estimate is

$$
\begin{equation*}
\widehat{f}_{n}(x)=e^{P_{x}(\widehat{a}, x)}=e^{\widehat{a}_{0}} . \tag{6.45}
\end{equation*}
$$

6.46 Remark. When $p=0, \widehat{f}_{n}$ reduces to kernel density estimation.

### 6.5 Multivariate Problems

Suppose now that the data are $d$-dimensional so that $X_{i}=\left(X_{i 1}, \ldots, X_{i d}\right)$. As we discussed in the two previous chapters, we can generalize the methods to higher dimensions quite easily in principle, though the curse of dimensionality implies that the accuracy of the estimator deteriorates quickly as dimension increases.

The kernel estimator can easily be generalized to $d$ dimensions. Most often, we use the product kernel

$$
\begin{equation*}
\widehat{f}_{n}(x)=\frac{1}{n h_{1} \cdots h_{d}} \sum_{i=1}^{n}\left\{\prod_{j=1}^{d} K\left(\frac{x_{j}-X_{i j}}{h_{j}}\right)\right\} . \tag{6.47}
\end{equation*}
$$

The risk is given by

$$
\begin{equation*}
R \approx \frac{1}{4} \sigma_{K}^{4}\left[\sum_{j=1}^{d} h_{j}^{4} \int f_{j j}^{2}(x) d x+\sum_{j \neq k} h_{j}^{2} h_{k}^{2} \int f_{j j} f_{k k} d x\right]+\frac{\left(\int K^{2}(x) d x\right)^{d}}{n h_{1} \cdots h_{d}} \tag{6.48}
\end{equation*}
$$

where $f_{j j}$ is the second partial derivative of $f$. The optimal bandwidth satisfies $h_{i}=O\left(n^{-1 /(4+d)}\right)$ leading to a risk of order $R=O\left(n^{-4 /(4+d)}\right)$. Again, we see that the risk increases quickly with dimension. To get a sense of how serious this problem is, consider the following table from Silverman (1986) which shows the sample size required to ensure a relative mean squared error less than 0.1 at 0 when the density is multivariate normal and the optimal bandwidth is selected.

| Dimension | Sample Size |
| ---: | ---: |
| 1 | 4 |
| 2 | 19 |
| 3 | 67 |
| 4 | 223 |
| 5 | 768 |
| 6 | 2790 |
| 7 | 10,700 |
| 8 | 43,700 |
| 9 | 187,000 |
| 10 | 842,000 |

This is bad news. If you do attempt to estimate a density in a high dimensional problem, you should not report the results without reporting confidence bands. The confidence band method described in Section 6.6 can be extended for the multivariate case, although we do not report the details here. These bands get very wide as $d$ increases. The problem is not the method of estimation; rather, the wide bands correctly reflect the difficulty of the problem.

### 6.6 Converting Density Estimation Into Regression

There is a useful trick for converting a density estimation problem into a regression problem. Then we can use all the regression methods from the previous chapter. This trick is an old idea but was recently made rigorous by Nussbaum (1996a) and Brown et al. (2005). By converting to regression, we can use all the tools we developed in the previous chapter, including the method for constructing confidence bands.

Suppose $X_{1}, \ldots, X_{n} \sim F$ with density $f=F^{\prime}$. For simplicity, suppose the data are on $[0,1]$. Divide the interval $[0,1]$ into $k$ equal width bins where $k \approx n / 10$. Define

$$
\begin{equation*}
Y_{j}=\sqrt{\frac{k}{n}} \times \sqrt{N_{j}+\frac{1}{4}} \tag{6.49}
\end{equation*}
$$

where $N_{j}$ is the number of observations in bin $j$. Then,

$$
\begin{equation*}
Y_{j} \approx r\left(t_{j}\right)+\sigma \epsilon_{j} \tag{6.50}
\end{equation*}
$$

where $\epsilon_{j} \sim N(0,1), \sigma=\sqrt{\frac{k}{4 n}}, r(x)=\sqrt{f(x)}$ and $t_{j}$ is the midpoint of the $j^{\text {th }}$ bin. To see why, let $B_{j}$ denote the $j^{\text {th }}$ bin and note that

$$
N_{j} \approx \operatorname{Poisson}\left(n \int_{B_{j}} f(x) d x\right) \approx \operatorname{Poisson}\left(\frac{n f\left(t_{j}\right)}{k}\right)
$$

so that $\mathbb{E}\left(N_{j}\right)=V\left(N_{j}\right) \approx n f\left(t_{j}\right) / k$. Applying the delta method, we see that $\mathbb{E}\left(Y_{j}\right) \approx \sqrt{f\left(t_{j}\right)}$ and $\mathbb{V}\left(Y_{j}\right) \approx k /(4 n)$.

We have thus converted the density estimation problem into a nonparametric regression problem with equally spaced $x_{i}$ s and constant variance. We can now apply any nonparametric regression method to get an estimate $\widehat{r}_{n}$ and take

$$
\widehat{f}_{n}(x)=\frac{\left(r^{+}(x)\right)^{2}}{\int_{0}^{1}\left(r^{+}(s)\right)^{2} d s}
$$

where $r^{+}(x)=\max \left\{\widehat{r}_{n}(x), 0\right\}$. In particular, we can construct confidence bands as in Chapter 5. It is important to note that binning is not a smoothing step; binning is being used to turn density estimation into regression.
6.51 Example. Figure 6.6 shows the method applied to data from the Bart Simpson distribution. The top plot shows the cross-validation score. The bottom plot shows the estimated density and 95 percent confidence bands.

### 6.7 Bibliographic Remarks

Kernel smoothing was invented by Rosenblatt (1956) and Parzen (1962). The cross-validation method is due to Rudemo (1982). Two very good books on density estimation are Scott (1992) and Silverman (1986). For information on a different approach called the scale-space approach see Chaudhuri and Marron (1999) and Chaudhuri and Marron (2000).

### 6.8 Appendix

proof of Theorem 6.11. For any $x, u \in B_{j}$,

$$
f(u)=f(x)+(u-x) f^{\prime}(x)+\frac{(u-x)^{2}}{2} f^{\prime \prime}(\widetilde{x})
$$

for some $\widetilde{x}$ between $x$ and $u$.
Hence,

$$
\begin{aligned}
p_{j}=\int_{B_{j}} f(u) d u & =\int_{B_{j}}\left(f(x)+(u-x) f^{\prime}(x)+\frac{(u-x)^{2}}{2} f^{\prime \prime}(\widetilde{x})\right) d u \\
& =f(x) h+h f^{\prime}(x)\left(h\left(j-\frac{1}{2}\right)-x\right)+O\left(h^{3}\right)
\end{aligned}
$$



FIGURE 6.6. Density estimation by regression. The data were binned and kernel regression was used on the square root of the counts. The top plot shows the cross-validation score by effective degrees of freedom. The bottom plot shows the estimate and a 95 percent confidence envelope. See Example 6.51.

Therefore, the bias $b(x)$ is

$$
\begin{aligned}
b(x) & =\mathbb{E}\left(\widehat{f}_{n}(x)\right)-f(x)=\frac{p_{j}}{h}-f(x) \\
& =\frac{f(x) h+h f^{\prime}(x)\left(h\left(j-\frac{1}{2}\right)-x\right)+O\left(h^{3}\right)}{h}-f(x) \\
& =f^{\prime}(x)\left(h\left(j-\frac{1}{2}\right)-x\right)+O\left(h^{2}\right) .
\end{aligned}
$$

By the mean value theorem we have, for some $\widetilde{x}_{j} \in B_{j}$, that

$$
\begin{aligned}
\int_{B_{j}} b^{2}(x) d x & =\int_{B_{j}}\left(f^{\prime}(x)\right)^{2}\left(h\left(j-\frac{1}{2}\right)-x\right)^{2} d x+O\left(h^{4}\right) \\
& =\left(f^{\prime}\left(\widetilde{x}_{j}\right)\right)^{2} \int_{B_{j}}\left(h\left(j-\frac{1}{2}\right)-x\right)^{2} d x+O\left(h^{4}\right)
\end{aligned}
$$

$$
=\left(f^{\prime}\left(\widetilde{x}_{j}\right)\right)^{2} \frac{h^{3}}{12}+O\left(h^{4}\right)
$$

Therefore,

$$
\begin{aligned}
\int_{0}^{1} b^{2}(x) d x & =\sum_{j=1}^{m} \int_{B_{j}} b^{2}(x) d x+O\left(h^{3}\right) \\
& =\sum_{j=1}^{m}\left(f^{\prime}\left(\widetilde{x}_{j}\right)\right)^{2} \frac{h^{3}}{12}+O\left(h^{3}\right) \\
& =\frac{h^{2}}{12} \sum_{j=1}^{m} h\left(f^{\prime}\left(\widetilde{x}_{j}\right)\right)^{2}+O\left(h^{3}\right) \\
& =\frac{h^{2}}{12} \int_{0}^{1}\left(f^{\prime}(x)\right)^{2} d x+o\left(h^{2}\right)
\end{aligned}
$$

Now consider the variance. By the mean value theorem, $p_{j}=\int_{B_{j}} f(x) d x=$ $h f\left(x_{j}\right)$ for some $x_{j} \in B_{j}$. Hence, with $v(x)=\mathbb{V}\left(\widehat{f}_{n}(x)\right.$,

$$
\begin{aligned}
\int_{0}^{1} v(x) d x & =\sum_{j=1}^{m} \int_{B_{j}} v(x) d x=\sum_{j=1}^{m} \int_{B_{j}} \frac{p_{j}\left(1-p_{j}\right)}{n h^{2}} \\
& =\frac{1}{n h^{2}} \sum_{j=1}^{m} \int_{B_{j}} p_{j}-\frac{1}{n h^{2}} \sum_{j=1}^{m} \int_{B_{j}} p_{j}^{2}=\frac{1}{n h}-\frac{1}{n h} \sum_{j=1}^{m} p_{j}^{2} \\
& =\frac{1}{n h}-\frac{1}{n h} \sum_{j=1}^{m} h^{2} f^{2}\left(x_{j}\right)=\frac{1}{n h}-\frac{1}{n} \sum_{j=1}^{m} h f^{2}\left(x_{j}\right) \\
& =\frac{1}{n h}-\frac{1}{n}\left(\int_{0}^{1} f^{2}(x) d x+o(1)\right)=\frac{1}{n h}+o\left(\frac{1}{n}\right)
\end{aligned}
$$

### 6.9 Exercises

1. Prove Theorem 6.27.
2. Let $X_{1}, \ldots, X_{n} \sim f$ and let $\widehat{f}_{n}$ be the kernel density estimator using the boxcar kernel:

$$
K(x)= \begin{cases}1 & -\frac{1}{2}<x<\frac{1}{2} \\ 0 & \text { otherwise }\end{cases}
$$

(a) Show that

$$
\mathbb{E}(\widehat{f}(x))=\frac{1}{h} \int_{x-(h / 2)}^{x+(h / 2)} f(y) d y
$$

and

$$
\mathbb{V}(\widehat{f}(x))=\frac{1}{n h^{2}}\left[\int_{x-(h / 2)}^{x+(h / 2)} f(y) d y-\left(\int_{x-(h / 2)}^{x+(h / 2)} f(y) d y\right)^{2}\right]
$$

(b) Show that if $h \rightarrow 0$ and $n h \rightarrow \infty$ as $n \rightarrow \infty$ then $\widehat{f}_{n}(x) \xrightarrow{\mathrm{P}} f(x)$.
3. Prove that $\widehat{J}(h)$ is an unbiased estimate of $J(h)$ for histograms and kernel density estimates.
4. Prove equation 6.35 .
5. Data on the salaries of the chief executive officer of 60 companies is available at:
http://lib.stat.cmu.edu/DASL/Datafiles/ceodat.html
Investigate the distribution of salaries using a histogram and a kernel density estimator. Use least squares cross-validation to choose the amount of smoothing. Also consider the Normal reference rule for picking a bandwidth for the kernel. There appear to be a few bumps in the density. Are they real? Use confidence bands to address this question. Finally, try using various kernel shapes and comment on the resulting estimates.
6. Get the data on fragments of glass collected in forensic work from the book website. Estimate the density of the first variable (refractive index) using a histogram and a kernel density estimator. Use cross-validation to choose the amount of smoothing. Experiment with different binwidths and bandwidths. Comment on the similarities and differences. Construct 95 percent confidence bands for your estimators. For the kernel, experiment with different kernel shapes.
7. Consider the data in Exercise 6. Examine the fit as a function of the bandwidth $h$. Do this by plotting the fit for many values of $h$. Add confidence bands to all the fits. If you are feeling very ambitious, read Chaudhuri and Marron (1999) and apply that method.
8. Prove that local likelihood density estimation reduces to kernel density estimation when the degree of the polynomial $p=0$.
9. Apply local polynomial density estimation to the data in Exercise 6.
10. Generate data from the Bart Simpson distribution (6.2). Compare kernel density estimation to the method in Section 6.6. Try the following sample sizes: $n=25,50,100,1,000$.

## 7

## Normal Means and Minimax Theory

In this chapter we will discuss the many Normal means problem which unifies some nonparametric problems and will be the basis for the methods in the next two chapters. The material in this chapter is more theoretical than in the rest of the book. If you are not interested in the theoretical details, I recommend reading sections $7.1,7.2$, and 7.3 and then skipping to the next chapter, referring back as needed. If you want more details on this topic, I recommend Johnstone (2003).

### 7.1 The Normal Means Model

Let $Z^{n}=\left(Z_{1}, \ldots, Z_{n}\right)$ where

$$
\begin{equation*}
Z_{i}=\theta_{i}+\sigma_{n} \epsilon_{i}, \quad i=1, \ldots, n \tag{7.1}
\end{equation*}
$$

$\epsilon_{1}, \ldots, \epsilon_{n}$ are independent, $\operatorname{Normal}(0,1)$ random variables,

$$
\theta^{n}=\left(\theta_{1}, \ldots, \theta_{n}\right) \in \mathbb{R}^{n}
$$

is a vector of unknown parameters and $\sigma_{n}$ is assumed known. Typically $\sigma_{n}=$ $\sigma / \sqrt{n}$ but we shall not assume this unless specifically noted. Sometimes we write $Z^{n}$ and $\theta^{n}$ as $Z$ and $\theta$. The model may appear to be parametric but the number of parameters is increasing at the same rate as the number of

| $\theta_{1}$ | $\theta_{2}$ | $\ldots$ | $\theta_{i}$ | $\ldots$ | $\theta_{n}$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $X_{11}$ | $X_{21}$ | $\cdots$ | $X_{i 1}$ | $\ldots$ | $X_{n 1}$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ |
| $X_{1 j}$ | $X_{2 j}$ | $\ldots$ | $X_{i j}$ | $\ldots$ | $X_{n j}$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ |
| $X_{1 n}$ | $X_{2 n}$ | $\ldots$ | $X_{i n}$ | $\vdots$ | $X_{n n}$ |
| $Z_{1}$ | $Z_{2}$ | $\ldots$ | $Z_{i}$ | $\cdots$ | $Z_{n}$ |

FIGURE 7.1. The Normal means model. $X_{i j}=\theta_{i}+N\left(0, \sigma^{2}\right)$ and $Z_{i}=n^{-1} \sum_{j=1}^{n} X_{i j}=\theta_{i}+\sigma_{n} \epsilon_{i}$ where $\sigma_{n}=\sigma / \sqrt{n}$. Estimating the parameters $\theta_{1}, \ldots, \theta_{n}$ from the $n$ column means $Z_{1}, \ldots, Z_{n}$ leads to the model (7.1) with $\sigma_{n}=\sigma / \sqrt{n}$.
data points. This model carries with it all the complexities and subtleties of a nonparametric problem. We will also consider an infinite-dimensional version of the model:

$$
\begin{equation*}
Z_{i}=\theta_{i}+\sigma_{n} \epsilon_{i}, \quad i=1,2, \ldots, \tag{7.2}
\end{equation*}
$$

where now the unknown parameter is $\theta=\left(\theta_{1}, \theta_{2}, \ldots\right)$.
Throughout this chapter we take $\sigma_{n}^{2}$ as known. In practice, we would need to estimate the variance using one of the methods discussed in Chapter 5. In this case, the exact results that follow may no longer hold but, under appropriate smoothness conditions, asymptotic versions of the results will hold.
7.3 Example. To provide some intuition for this model, suppose that we have data $X_{i j}=\theta_{i}+\sigma \delta_{i j}$ where $1 \leq i, j \leq n$ and the $\delta_{i j}$ are independent $\mathrm{N}(0,1)$ random variables. This is simply a one-way analysis of variance model; see Figure 7.1. Let $Z_{i}=n^{-1} \sum_{j=1}^{n} X_{i j}$. Then the model (7.1) holds with $\sigma_{n}=$ $\sigma / \sqrt{n}$. We get the infinite version (7.2) by having infinitely many columns in Figure 7.1 (but still $n$ rows).

Given an estimator $\widehat{\theta}^{n}=\left(\widehat{\theta}_{1}, \ldots, \widehat{\theta}_{n}\right)$ we will use the squared error loss

$$
L\left(\widehat{\theta}^{n}, \theta^{n}\right)=\sum_{i=1}^{n}\left(\widehat{\theta}_{i}-\theta_{i}\right)^{2}=\left\|\widehat{\theta}^{n}-\theta^{n}\right\|^{2}
$$

with risk function

$$
R\left(\widehat{\theta}^{n}, \theta^{n}\right)=\mathbb{E}_{\theta}\left(L\left(\widehat{\theta}^{n}, \theta^{n}\right)\right)=\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-\theta_{i}\right)^{2}
$$

An obvious choice for an estimator of $\theta^{n}$ is $\widehat{\theta}^{n}=Z^{n}$. This estimator has impressive credentials: it is the maximum likelihood estimator, it is the min-
imum variance unbiased estimator and it is the Bayes estimator under a flat prior. Nonetheless, it is a poor estimator. Its risk is

$$
R\left(Z^{n}, \theta^{n}\right)=\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(Z_{i}-\theta_{i}\right)^{2}=\sum_{i=1}^{n} \sigma_{n}^{2}=n \sigma_{n}^{2}
$$

We shall see that there are estimators with substantially smaller risk.
Before we explain how we can improve on the MLE, let's first see how the normal means problem relates to nonparametric regression and density estimation. To do that, we need to review some theory about function spaces.

### 7.2 Function Spaces

Let $L_{2}(a, b)$ denote the set of functions $f:[a, b] \rightarrow \mathbb{R}$ such that $\int_{a}^{b} f^{2}(x) d x<$ $\infty$. Unless otherwise indicated, assume that $a=0$ and $b=1$. The inner product between two functions $f$ and $g$ in $L_{2}(a, b)$ is $\int_{a}^{b} f(x) g(x) d x$ and the norm of $f$ is $\|f\|=\sqrt{\int_{a}^{b} f^{2}(x) d x}$. A sequence of functions $\phi_{1}, \phi_{2}, \ldots$ is called orthonormal if $\left\|\phi_{j}\right\|=1$ for all $j$ (normalized) and $\int_{a}^{b} \phi_{i}(x) \phi_{j}(x) d x=0$ for $i \neq j$ (orthogonal). The sequence is complete if the only function that is orthogonal to each $\phi_{j}$ is the zero function. A complete, orthonormal set of functions forms a basis, meaning that if $f \in L_{2}(a, b)$ then $f$ can be expanded in the basis:
7.4 Theorem. If $f \in L_{2}(a, b)$ then ${ }^{1}$

$$
\begin{equation*}
f(x)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(x) \tag{7.5}
\end{equation*}
$$

where

$$
\begin{equation*}
\theta_{j}=\int_{a}^{b} f(x) \phi_{j}(x) d x \tag{7.6}
\end{equation*}
$$

Furthermore,

$$
\begin{equation*}
\int_{a}^{b} f^{2}(x) d x=\sum_{j=1}^{\infty} \theta_{j}^{2} \tag{7.7}
\end{equation*}
$$

which is known as Parseval's identity.

[^12]An example of an orthonormal basis for $L_{2}(0,1)$ is the cosine basis

$$
\phi_{0}(x)=1, \quad \phi_{j}(x)=\sqrt{2} \cos (2 \pi j x), \quad j=1,2, \ldots .
$$

Another example is the Legendre basis defined on $(-1,1)$ :

$$
P_{0}(x)=1, \quad P_{1}(x)=x, \quad P_{2}(x)=\frac{1}{2}\left(3 x^{2}-1\right), \quad P_{3}(x)=\frac{1}{2}\left(5 x^{3}-3 x\right), \ldots
$$

These polynomials are defined by the relation

$$
P_{n}(x)=\frac{1}{2^{n} n!} \frac{d^{n}}{d x^{n}}\left(x^{2}-1\right)^{n}
$$

The Legendre polynomials are orthogonal but not orthonormal since

$$
\int_{-1}^{1} P_{n}^{2}(x) d x=\frac{2}{2 n+1} .
$$

However, we can define modified Legendre polynomials $Q_{n}(x)=\sqrt{(2 n+1) / 2}$ $P_{n}(x)$ which then form an orthonormal basis for $L_{2}(-1,1)$.

Next we introduce Sobolev spaces, which are sets of smooth functions. Let $D^{j} f$ denote the $j^{\text {th }}$ weak derivative ${ }^{2}$ of $f$.
7.8 Definition. The Sobolev space of order $m$, is defined by

$$
W(m)=\left\{f \in L_{2}(0,1): D^{m} f \in L_{2}(0,1)\right\} .
$$

The Sobolev space of order $m$ and radius $c$, is defined by

$$
W(m, c)=\left\{f: f \in W(m),\left\|D^{m} f\right\|^{2} \leq c^{2}\right\}
$$

The periodic Sobolev class is

$$
\widetilde{W}(m, c)=\left\{f \in W(m, c): D^{j} f(0)=D^{j} f(1), j=0, \ldots, m-1\right\} .
$$

An ellipsoid is a set of the form

$$
\begin{equation*}
\Theta=\left\{\theta: \sum_{j=1}^{\infty} a_{j}^{2} \theta_{j}^{2} \leq c^{2}\right\} \tag{7.9}
\end{equation*}
$$

where $a_{j}$ is a sequence of numbers such that $a_{j} \rightarrow \infty$ as $j \rightarrow \infty$.

[^13]7.10 Definition. If $\Theta$ is an ellipsoid and if $a_{j}^{2} \sim(\pi j)^{2 m}$ as $j \rightarrow \infty$, we call $\Theta a$ Sobolev ellipsoid or a Sobolev body and we denote it by $\Theta(m, c)$.

Now we relate Sobolev spaces to Sobolev ellipsoids.
7.11 Theorem. Let $\left\{\phi_{j}, j=0,1, \ldots\right\}$ be the Fourier basis:
$\phi_{1}(x)=1, \quad \phi_{2 j}(x)=\frac{1}{\sqrt{2}} \cos (2 j \pi x), \quad \phi_{2 j+1}(x)=\frac{1}{\sqrt{2}} \sin (2 j \pi x), \quad j=1,2, \ldots$.
Then,

$$
\begin{equation*}
\widetilde{W}(m, c)=\left\{f: f=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}, \sum_{j=1}^{\infty} a_{j}^{2} \theta_{j}^{2} \leq c^{2}\right\} \tag{7.12}
\end{equation*}
$$

where $a_{j}=(\pi j)^{m}$ for $j$ even and $a_{j}=(\pi(j-1))^{m}$ for $j$ odd.
Thus, a Sobolev space corresponds to a Sobolev ellipsoid with $a_{j} \sim(\pi j)^{2 m}$. It is also possible to relate the class $W(m, c)$ to an ellipsoid although the details are more complicated; see Nussbaum (1985).

In Sobolev spaces, smooth functions have small coefficients $\theta_{j}$ when $j$ is large, otherwise $\sum_{j} \theta_{j}^{2}(\pi j)^{2 m}$ will blow up. Thus, to smooth a function, we shrink the $\theta_{j}$ s closer to zero. Hence:
smoothing $f$ corresponds to shrinking the $\theta_{j}$ 's towards zero for large $j$.

A generalization of Sobolev spaces are Besov spaces. These include Sobolev spaces as a special case but they also include functions with less smoothness. We defer discussion of Besov spaces until Chapter 9.

### 7.3 Connection to Regression and Density Estimation

Consider the nonparametric regression model

$$
\begin{equation*}
Y_{i}=f(i / n)+\sigma \epsilon_{i}, \quad i=1, \ldots, n \tag{7.13}
\end{equation*}
$$

where $\epsilon_{i} \sim N(0,1), \sigma$ is known and $f \in L_{2}(0,1)$. Let $\phi_{1}, \phi_{2}, \ldots$ be an orthonormal basis and write $f(x)=\sum_{i=1}^{\infty} \theta_{j} \phi_{j}(x)$ where $\theta_{j}=\int f(x) \phi_{j}(x) d x$. First, approximate $f$ by the finite series $f(x) \approx \sum_{i=1}^{n} \theta_{j} \phi_{j}(x)$. Now define

$$
\begin{equation*}
Z_{j}=\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}(i / n) \tag{7.14}
\end{equation*}
$$

for $j=1, \ldots, n$. The random variable $Z_{j}$ has a Normal distribution since $Z_{j}$ is a linear combination of Normals. The mean of $Z_{j}$ is

$$
\begin{aligned}
\mathbb{E}\left(Z_{j}\right) & =\frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\left(Y_{i}\right) \phi_{j}(i / n)=\frac{1}{n} \sum_{i=1}^{n} f(i / n) \phi_{j}(i / n) \\
& \approx \int f(x) \phi_{j}(x) d x=\theta_{j} .
\end{aligned}
$$

The variance is

$$
\begin{aligned}
\mathbb{V}\left(Z_{j}\right) & =\frac{1}{n^{2}} \sum_{i=1}^{n} \mathbb{V}\left(Y_{i}\right) \phi_{j}^{2}(i / n)=\frac{\sigma^{2}}{n} \frac{1}{n} \sum_{i=1}^{n} \phi_{j}^{2}(i / n) \\
& \approx \frac{\sigma^{2}}{n} \int \phi_{j}^{2}(x) d x=\frac{\sigma^{2}}{n} \equiv \sigma_{n}^{2}
\end{aligned}
$$

A similar calculation shows that $\operatorname{Cov}\left(Z_{j}, Z_{k}\right) \approx 0$. We conclude that the $Z_{j}$ are approximately independent and

$$
\begin{equation*}
Z_{j} \sim N\left(\theta_{j}, \sigma_{n}^{2}\right), \quad \sigma_{n}^{2}=\frac{\sigma^{2}}{n} \tag{7.15}
\end{equation*}
$$

We have thus converted the problem of estimating $f$ into the problem of estimating the means of $n$ Normal random variables as in (7.1) with $\sigma_{n}^{2}=$ $\sigma^{2} / n$. Also, squared error loss for $f$ corresponds to squared error loss for $\theta$ since, by Parseval's identity, if $\widehat{f}_{n}(x)=\sum_{j=1}^{\infty} \widehat{\theta}_{j} \phi_{j}(x)$,

$$
\begin{equation*}
\left\|\widehat{f}_{n}-f\right\|^{2}=\int\left(\widehat{f}_{n}(x)-f(x)\right)^{2} d x=\sum_{j=1}^{\infty}\left(\widehat{\theta}_{j}-\theta_{j}\right)^{2}=\|\widehat{\theta}-\theta\|^{2} \tag{7.16}
\end{equation*}
$$

where $\|\theta\|=\sqrt{\sum_{j} \theta_{j}^{2}}$.
It turns out that other nonparametric problems, such as density estimation, can also be connected to the Normal means problem. In the case of density estimation, it is the square root of the density that appears in the white noise problem. In this sense, the many Normal means problem serves as a unifying framework for many nonparametric models. See Nussbaum (1996a), Claeskens and Hjort (2004) and the appendix for more details.

### 7.4 Stein's Unbiased Risk Estimator (SURE)

Let $\widehat{\theta}$ be an estimate of $\theta$. It will be useful to have an estimate of the risk of $\widehat{\theta}$. In previous chapters we used cross-validation to estimate risk. In the present context there is a more elegant method for risk estimation due to Stein (1981) known as Stein's unbiased risk estimator or SURE.
7.17 Theorem (Stein). Let $Z \sim N_{n}(\theta, V)$, let $\widehat{\theta}=\widehat{\theta}(Z)$ be an estimate of $\theta$ and let $g\left(Z_{1}, \ldots, Z_{n}\right)=\widehat{\theta}-Z$. Note that $g$ maps $\mathbb{R}^{n}$ to $\mathbb{R}^{n}$. Define

$$
\begin{equation*}
\widehat{R}(z)=\operatorname{tr}(V)+2 \operatorname{tr}(V D)+\sum_{i} g_{i}^{2}(z) \tag{7.18}
\end{equation*}
$$

where $\operatorname{tr}$ denotes the trace of a matrix, $g_{i}=\widehat{\theta}_{i}-Z_{i}$ and the $(i, j)$ component of $D$ is the partial derivative of the $i^{\text {th }}$ component of $g\left(z_{1}, \ldots, z_{n}\right)$ with respect to $z_{j}$. If $g$ is weakly differentiable ${ }^{3}$ then

$$
\mathbb{E}_{\theta}(\widehat{R}(Z))=R(\theta, \widehat{\theta})
$$

If we apply Theorem 7.17 to the model (7.1) we get the following.

## The sure Formula for the Normal Means Model

Let $\widehat{\theta}$ be a weakly differentiable estimator of $\theta$ in model (7.1). An unbiased estimate of the risk of $\widehat{\theta}$ is

$$
\begin{equation*}
\widehat{R}(z)=n \sigma_{n}^{2}+2 \sigma_{n}^{2} \sum_{i=1}^{n} D_{i}+\sum_{i=1}^{n} g_{i}^{2} \tag{7.19}
\end{equation*}
$$

where $g\left(Z_{1}, \ldots, Z_{n}\right)=\widehat{\theta}^{n}-Z^{n}$ and $D_{i}=\partial g\left(z_{1}, \ldots, z_{n}\right) / \partial z_{i}$.
Proof of Theorem 7.17. We will prove the case where $V=\sigma^{2} I$. If $X \sim N\left(\mu, \sigma^{2}\right)$ then $\mathbb{E}(g(X)(X-\mu))=\sigma^{2} \mathbb{E} g^{\prime}(X)$. (This is known as Stein's Lemma and it can be proved using integration by parts. See Exercise 4.) Hence, $\sigma^{2} \mathbb{E}_{\theta} D_{i}=\mathbb{E}_{\theta} g_{i}\left(Z_{i}-\theta\right)$ and

$$
\begin{aligned}
\mathbb{E}_{\theta}(\widehat{R}(Z))= & n \sigma^{2}+2 \sigma^{2} \sum_{i=1}^{n} \mathbb{E}_{\theta} D_{i}+\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-Z_{i}\right)^{2} \\
= & n \sigma^{2}+2 \sum_{i=1}^{n} \mathbb{E}_{\theta}\left(g_{i}\left(Z_{i}-\theta_{i}\right)\right)+\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-Z_{i}\right)^{2} \\
= & \sum_{i=1}^{n} \mathbb{E}_{\theta}\left(Z_{i}-\theta_{i}\right)^{2}+2 \sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\left(\widehat{\theta}_{i}-Z_{i}\right)\left(Z_{i}-\theta_{i}\right)\right) \\
& +\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-Z_{i}\right)^{2} \\
= & \sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-Z_{i}+Z_{i}-\theta_{i}\right)^{2}=\sum_{i=1}^{n} \mathbb{E}_{\theta}\left(\widehat{\theta}_{i}-\theta_{i}\right)^{2}=R(\widehat{\theta}, \theta)
\end{aligned}
$$

[^14]7.20 Example. Let $V=\sigma^{2} I$. Consider $\widehat{\theta}=Z$. Then $g(z)=(0, \ldots, 0)$ and $\widehat{R}(Z)=n \sigma^{2}$. In this case, $\widehat{R}$ is equal to the true risk. Now consider the linear estimator $\widehat{\theta}=b Z=\left(b Z_{1}, \ldots, b Z_{n}\right)$. Hence, $g(Z)=b Z-Z=(b-1) Z$ and $D_{i}=b-1$. Therefore, $\widehat{R}(Z)=(2 b-1) n \sigma^{2}+(1-b)^{2} \sum_{i=1}^{n} Z_{i}^{2}$. Next consider the soft threshold estimator defined by
\[

\widehat{\theta}_{i}= $$
\begin{cases}Z_{i}+\lambda & Z_{i}<-\lambda  \tag{7.21}\\ 0 & -\lambda \leq Z_{i} \leq \lambda \\ Z_{i}-\lambda & Z_{i}>\lambda\end{cases}
$$
\]

where $\lambda>0$ is a constant. We can write this estimator more succinctly as

$$
\widehat{\theta}_{i}=\operatorname{sign}\left(Z_{i}\right)\left(\left|Z_{i}\right|-\lambda\right)_{+} .
$$

In Exercise 5 you will show that the SURE formula gives

$$
\begin{equation*}
\widehat{R}(Z)=\sum_{i=1}^{n}\left(\sigma^{2}-2 \sigma^{2} I\left(\left|Z_{i}\right| \leq \lambda\right)+\min \left(Z_{i}^{2}, \lambda^{2}\right)\right) \tag{7.22}
\end{equation*}
$$

Finally, consider the hard threshold estimator defined by

$$
\widehat{\theta}_{i}= \begin{cases}Z_{i} & \left|Z_{i}\right|>\lambda  \tag{7.23}\\ 0 & \left|Z_{i}\right| \leq \lambda\end{cases}
$$

where $\lambda>0$ is a constant. It is tempting to use sure but this is inappropriate because this estimator is not weakly differentiable.
7.24 Example (Model selection). For each $S \subset\{1, \ldots, n\}$ define

$$
\begin{equation*}
\widehat{\theta}_{S}=Z_{i} I(i \in S) \tag{7.25}
\end{equation*}
$$

We can think of $S$ as a submodel which says that $Z_{i} \sim N\left(\theta_{i}, \sigma_{n}^{2}\right)$ for $i \in S$ and $Z_{i} \sim N\left(0, \sigma_{n}^{2}\right)$ for $i \notin S$. Then $\widehat{\theta}_{S}$ is the estimator of $\theta$ assuming the model $S$. The true risk of $\hat{\theta}_{S}$ is

$$
R\left(\widehat{\theta}_{S}, \theta\right)=\sigma_{n}^{2}|S|+\sum_{i \in S^{c}} \theta_{i}^{2}
$$

where $|S|$ denotes the number of points in $S$. Replacing $\theta_{i}^{2}$ in the risk with its unbiased estimator $Z_{i}^{2}-\sigma_{n}^{2}$ yields the risk estimator

$$
\begin{equation*}
\widehat{R}_{S}=\sigma_{n}^{2}|S|+\sum_{i \in S^{c}}\left(Z_{i}^{2}-\sigma_{n}^{2}\right) \tag{7.26}
\end{equation*}
$$

It is easy to check that this corresponds to the SURE formula. Now let $\mathcal{S}$ be some class of sets where each $S \in \mathcal{S}$ is a subset of $\{1, \ldots, n\}$. Choosing $S \in \mathcal{S}$ to minimize $\widehat{R}_{S}$ is an example of model selection. The special case where

$$
\mathcal{S}=\{\emptyset,\{1\},\{1,2\}, \ldots,\{1,2, \ldots, n\}\}
$$

is called nested subset selection. Taking $\mathcal{S}$ to be all subsets of $\{1, \ldots, n\}$ corresponds to all possible subsets. For any fixed model $S$, we expect that $\widehat{R}_{S}$ will be close to $R\left(\widehat{\theta}_{S}, \theta\right)$. However, this does not guarantee that $\widehat{R}_{S}$ is close to $R\left(\widehat{\theta}_{S}, \theta\right)$ uniformly over $\mathcal{S}$. See Exercise 10.

### 7.5 Minimax Risk and Pinsker's Theorem

If $\Theta_{n}$ is a subset of $\mathbb{R}^{n}$, we define the minimax risk over $\Theta_{n}$ by

$$
\begin{equation*}
R_{n} \equiv R\left(\Theta_{n}\right) \equiv \inf _{\widehat{\theta}} \sup _{\theta \in \Theta_{n}} R(\widehat{\theta}, \theta) \tag{7.27}
\end{equation*}
$$

where the infimum is over all estimators. Two questions we will address are: (i) what is the value of the minimax risk $R\left(\Theta_{n}\right)$ ? and (ii) can we find an estimator that achieves this risk?

The following theorem ${ }^{4}$ gives the exact, limiting form of the minimax risk for the $L_{2}$ ball

$$
\Theta_{n}(c)=\left\{\left(\theta_{1}, \ldots, \theta_{n}\right): \sum_{i=1}^{n} \theta_{i}^{2} \leq c^{2}\right\}
$$

7.28 Theorem (Pinsker's theorem). Assume the model (7.1) with $\sigma_{n}^{2}=\sigma^{2} / n$. For any $c>0$,

$$
\begin{equation*}
\liminf _{n \rightarrow \infty} \inf _{\widehat{\theta}} \sup _{\theta \in \Theta_{n}(c)} R(\widehat{\theta}, \theta)=\frac{\sigma^{2} c^{2}}{\sigma^{2}+c^{2}} \tag{7.29}
\end{equation*}
$$

The right-hand side of (7.29) gives an exact expression for the (asymptotic) minimax risk. This expression is strictly smaller than $\sigma^{2}$ which is the risk for the maximum likelihood estimator. Later, we will introduce the James-Stein estimator which asymptotically achieves this risk. The proof of the theorem, which is in the appendix, is a bit technical and may be skipped without loss of continuity. Here is the basic idea behind the proof.

First, we note that the estimator with coordinates $\widehat{\theta}_{j}=c^{2} Z_{j} /\left(\sigma^{2}+c^{2}\right)$ has risk bounded above by $\sigma^{2} c^{2} /\left(\sigma^{2}+c^{2}\right)$. Hence,

$$
\begin{equation*}
R_{n} \leq \frac{\sigma^{2} c^{2}}{\sigma^{2}+c^{2}} \tag{7.30}
\end{equation*}
$$

If we could find a prior $\pi$ on $\Theta_{n}(c)$ whose posterior mean $\widetilde{\theta}$ also has risk $\sigma^{2} c^{2} /\left(\sigma^{2}+c^{2}\right)$ then we could argue that, for any estimator $\hat{\theta}$, we have

$$
\begin{equation*}
\frac{\sigma^{2} c^{2}}{\sigma^{2}+c^{2}}=\int R(\theta, \widetilde{\theta}) d \pi(\theta) \leq \int R(\theta, \widehat{\theta}) d \pi(\theta) \leq \sup _{\theta \in \Theta_{n}} R(\theta, \widehat{\theta})=R_{n} \tag{7.31}
\end{equation*}
$$

[^15]Combining (7.30) and (7.31) would yield $R_{n}=\sigma^{2} c^{2} /\left(\sigma^{2}+c^{2}\right)$. The proof is essentially an approximate version of this argument. One finds a prior over all of $\mathbb{R}^{n}$ whose risk is arbitrarily close to $\sigma^{2} c^{2} /\left(\sigma^{2}+c^{2}\right)$ and one then shows that the prior asymptotically concentrates on $\Theta_{n}(c)$.

Now let us see how minimax theory works for smooth functions.
7.32 Theorem (Pinsker's theorem for Sobolev ellipsoids). Let

$$
\begin{equation*}
Z_{j}=\theta_{j}+\frac{\sigma}{\sqrt{n}} \epsilon_{j}, \quad j=1,2, \ldots \tag{7.33}
\end{equation*}
$$

where $\epsilon_{1}, \epsilon_{2}, \ldots \sim N(0,1)$. Assume that $\theta \in \Theta(m, c)$, a Sobolev ellipsoid (recall Definition 7.10). Let $R_{n}$ denote the minimax risk over $\Theta(m, c)$. Then,

$$
\begin{equation*}
\lim _{n \rightarrow \infty} n^{2 m /(2 m+1)} R_{n}=\left(\frac{\sigma}{\pi}\right)^{2 m /(2 m+1)} c^{2 /(2 m+1)} P_{m} \tag{7.34}
\end{equation*}
$$

where

$$
\begin{equation*}
P_{m}=\left(\frac{m}{m+1}\right)^{2 m /(2 m+1)}(2 m+1)^{1 /(2 m+1)} \tag{7.35}
\end{equation*}
$$

is the Pinsker constant. Hence, the minimax rate is $n^{-2 m /(2 m+1)}$, that is,

$$
0<\lim _{n \rightarrow \infty} n^{2 m /(2 m+1)} R_{n}<\infty
$$

Here is a more general version of the theorem.
7.36 Theorem (Pinsker's theorem for ellipsoids). Let

$$
\Theta=\left\{\theta: \sum_{j=1}^{\infty} a_{j} \theta_{j}^{2} \leq c^{2}\right\} .
$$

The set $\Theta$ is called an ellipsoid. Assume that $a_{j} \rightarrow \infty$ as $j \rightarrow \infty$. Let

$$
R_{n}=\inf _{\widehat{\theta}} \sup _{\theta \in \Theta} R(\widehat{\theta}, \theta)
$$

denote the minimax risk and let

$$
R_{n}^{L}=\inf _{\hat{\theta} \in \mathcal{L}} \sup _{\theta \in \Theta} R(\widehat{\theta}, \theta)
$$

denote the minimax linear risk where $\mathcal{L}$ is the set of linear estimators of the form $\widehat{\theta}=\left(w_{1} Z_{1}, w_{2} Z_{2}, \ldots\right)$. Then:
(1) linear estimators are asymptotically minimax: $R_{n} \sim R_{n}^{L}$ as $n \rightarrow \infty$;
(2) the minimax linear risk satisfies

$$
R_{n}^{L}=\frac{\sigma^{2}}{n} \sum_{i}\left(1-\frac{a_{i}}{\mu}\right)_{+}
$$

where $\mu$ solves

$$
\frac{\sigma^{2}}{n} \sum_{i} a_{i}\left(\mu-a_{i}\right)_{+}=c^{2}
$$

(3) The linear minimax estimator is $\widehat{\theta}_{i}=w_{i} Z_{i}$ where $w_{i}=\left[1-\left(a_{i} / \mu\right)\right]_{+}$.
(4) The linear minimax estimator is Bayes ${ }^{5}$ for the prior with independent components such that $\theta_{i} \sim N\left(0, \tau_{i}^{2}\right), \tau_{i}^{2}=\left(\sigma^{2} / n\right)\left(\mu / a_{i}-1\right)_{+}$.

### 7.6 Linear Shrinkage and the James-Stein Estimator

Let us now return to model (7.1) and see how we can improve on the MLE using linear estimators. A linear estimator is an estimator of the form $\widehat{\theta}=b Z=$ $\left(b Z_{1}, \ldots, b Z_{n}\right)$ where $0 \leq b \leq 1$. Linear estimators are shrinkage estimators since they shrink $Z$ towards the origin. We denote the set of linear shrinkage estimators by $\mathcal{L}=\{b Z: b \in[0,1]\}$.

The risk of a linear estimator is easy to compute. From the basic biasvariance breakdown we have

$$
\begin{equation*}
R(b Z, \theta)=(1-b)^{2}\|\theta\|_{n}^{2}+n b^{2} \sigma_{n}^{2} \tag{7.37}
\end{equation*}
$$

where $\|\theta\|_{n}^{2}=\sum_{i=1}^{n} \theta_{i}^{2}$. The risk is minimized by taking

$$
b_{*}=\frac{\|\theta\|_{n}^{2}}{n \sigma_{n}^{2}+\|\theta\|_{n}^{2}} .
$$

We call $b_{*} Z$ the ideal linear estimator. The risk of this ideal linear estimator is

$$
\begin{equation*}
R\left(b_{*} Z, \theta\right)=\frac{n \sigma_{n}^{2}\|\theta\|_{n}^{2}}{n \sigma_{n}^{2}+\|\theta\|_{n}^{2}} \tag{7.38}
\end{equation*}
$$

Thus we have proved:

### 7.39 Theorem.

$$
\begin{equation*}
\inf _{\hat{\theta} \in \mathcal{L}} R(\widehat{\theta}, \theta)=\frac{n \sigma_{n}^{2}\|\theta\|_{n}^{2}}{n \sigma_{n}^{2}+\|\theta\|_{n}^{2}} \tag{7.40}
\end{equation*}
$$

We can't use the estimator $b_{*} Z$ because $b_{*}$ depends on the unknown parameter $\theta$. For this reason we call $R\left(b_{*} Z, \theta\right)$ the linear oracular risk since the risk could only be obtained by an "oracle" that knows $\|\theta\|_{n}^{2}$. We shall now show that the James-Stein estimator nearly achieves the risk of the ideal oracle.

[^16]The James-Stein estimator of $\theta$ is defined by

$$
\begin{equation*}
\widehat{\theta}^{J S}=\left(1-\frac{(n-2) \sigma_{n}^{2}}{\sum_{i=1}^{n} Z_{i}^{2}}\right) Z \tag{7.41}
\end{equation*}
$$

We'll see in Theorem 7.48 that this estimator is asymptotically optimal.
7.42 Theorem. The risk of the James-Stein estimator satisfies the following bound:

$$
\begin{equation*}
R\left(\widehat{\theta}^{J S}, \theta\right) \leq 2 \sigma_{n}^{2}+\frac{(n-2) \sigma_{n}^{2}\|\theta\|_{n}^{2}}{(n-2) \sigma_{n}^{2}+\|\theta\|_{n}^{2}} \leq 2 \sigma_{n}^{2}+\frac{n \sigma_{n}^{2}\|\theta\|_{n}^{2}}{n \sigma_{n}^{2}+\|\theta\|_{n}^{2}} \tag{7.43}
\end{equation*}
$$

where $\|\theta\|_{n}^{2}=\sum_{i=1}^{n} \theta_{i}^{2}$.
Proof. Write $\widehat{\theta}^{J S}=Z+g(Z)$ where $g(z)=-(n-2) \sigma_{n}^{2} z / \sum_{i} z_{i}^{2}$. Hence

$$
D_{i}=\frac{\partial g_{i}}{\partial z_{i}}=-(n-2) \sigma_{n}^{2}\left(\frac{1}{\sum_{i} z_{i}^{2}}-\frac{2 z_{i}^{2}}{\left(\sum_{i} z_{i}^{2}\right)^{2}}\right)
$$

and

$$
\sum_{i=1}^{n} D_{i}=-\frac{(n-2)^{2} \sigma_{n}^{2}}{\sum_{i=1}^{n} z_{i}^{2}}
$$

Plugging this into the SURE formula (7.19) yields

$$
\widehat{R}(Z)=n \sigma_{n}^{2}-\frac{(n-2)^{2} \sigma_{n}^{4}}{\sum_{i} Z_{i}^{2}}
$$

Hence, the risk is

$$
\begin{equation*}
R\left(\widehat{\theta}^{J S}, \theta\right)=\mathbb{E}(\widehat{R}(Z))=n \sigma_{n}^{2}-(n-2)^{2} \sigma_{n}^{4} \mathbb{E}\left(\frac{1}{\sum_{i} Z_{i}^{2}}\right) \tag{7.44}
\end{equation*}
$$

Now $Z_{i}^{2}=\sigma_{n}^{2}\left(\left(\theta_{i} / \sigma_{n}\right)+\epsilon_{i}\right)^{2}$ and so $\sum_{i=1}^{n} Z_{i}^{2} \sim \sigma_{n}^{2} W$ where $W$ is noncentral $\chi^{2}$ with $n$ degrees of freedom and noncentrality parameter $\delta=\sum_{i=1}^{n}\left(\theta_{i}^{2} / \sigma_{n}^{2}\right)$. Using a well-known result about noncentral $\chi^{2}$ random variables, we can then write $W \sim \chi_{n+2 K}^{2}$ where $K \sim \operatorname{Poisson}(\delta / 2)$. Recall that (for $\left.n>2\right) \mathbb{E}\left(1 / \chi_{n}^{2}\right)=$ $1 /(n-2)$. So,

$$
\begin{aligned}
\mathbb{E}_{\theta}\left[\frac{1}{\sum_{i} Z_{i}^{2}}\right] & =\left(\frac{1}{\sigma_{n}^{2}}\right) \mathbb{E}\left[\frac{1}{\chi_{n+2 K}^{2}}\right]=\left(\frac{1}{\sigma_{n}^{2}}\right) \mathbb{E}\left(E\left[\left.\frac{1}{\chi_{n+2 K}^{2}} \right\rvert\, K\right]\right) \\
& =\left(\frac{1}{\sigma_{n}^{2}}\right) \mathbb{E}\left[\frac{1}{n-2+2 K}\right] \\
& \geq\left(\frac{1}{\sigma_{n}^{2}}\right) \frac{1}{(n-2)+\sigma_{n}^{-2} \sum_{i=1}^{n} \theta_{i}^{2}} \quad \text { from Jensen's inequality } \\
& =\frac{1}{(n-2) \sigma_{n}^{2}+\sum_{i=1}^{n} \theta_{i}^{2}}
\end{aligned}
$$

Substituting into (7.44) we get the first inequality. The second inequality follows from simple algebra.
7.45 Remark. The modified James-Stein estimator is defined by

$$
\begin{equation*}
\widehat{\theta}=\left(1-\frac{n \sigma_{n}^{2}}{\sum_{i} Z_{i}^{2}}\right)_{+} Z \tag{7.46}
\end{equation*}
$$

where $(a)_{+}=\max \{a, 0\}$. The change from $n-2$ to $n$ leads to a simpler expression and for large $n$ this has negligible effect. Taking the positive part of the shrinkage factor cannot increase the risk. In practice, the modified James-Stein estimator is often preferred.

The next result shows that the James-Stein estimator nearly achieves the risk of the linear oracle.
7.47 Theorem (James-Stein oracle inequality). Let $\mathcal{L}=\{b Z: b \in \mathbb{R}\}$ denote the class of linear estimators. For all $\theta \in \mathbb{R}^{n}$,

$$
\inf _{\hat{\theta} \in \mathcal{L}} R(\widehat{\theta}, \theta) \leq R\left(\widehat{\theta}^{J S}, \theta\right) \leq 2 \sigma_{n}^{2}+\inf _{\hat{\theta} \in \mathcal{L}} R(\widehat{\theta}, \theta)
$$

Proof. This follows from (7.38) and Theorem 7.42.
Here is another perspective on the James-Stein estimator. Let $\widehat{\theta}=b Z$. Stein's unbiased risk estimator is $\widehat{R}(Z)=n \sigma_{n}^{2}+2 n \sigma_{n}^{2}(b-1)+(b-1)^{2} \sum_{i=1}^{n} Z_{i}^{2}$ which is minimized at

$$
\widehat{b}=1-\frac{n \sigma_{n}^{2}}{\sum_{i=1}^{n} Z_{i}^{2}}
$$

yielding the estimator

$$
\widehat{\theta}=\widehat{b} Z=\left(1-\frac{n \sigma_{n}^{2}}{\sum_{i=1}^{n} Z_{i}^{2}}\right) Z
$$

which is essentially the James-Stein estimator.
We can now show that the James-Stein estimator achieves the Pinsker bound (7.29) and so is asymptotically minimax.
7.48 Theorem. Let $\sigma_{n}^{2}=\sigma^{2} / n$. The James-Stein estimator is asymptotically minimax, that is,

$$
\lim _{n \rightarrow \infty} \sup _{\theta \in \Theta_{n}(c)} R\left(\widehat{\theta}^{J S}, \theta\right)=\frac{\sigma^{2} c^{2}}{\sigma^{2}+c^{2}}
$$

Proof. Follows from Theorem 7.42 and 7.28 .
7.49 Remark. The James-Stein estimator is adaptive in the sense that it achieves the minimax bound over $\Theta_{n}(c)$ without knowledge of the parameter $c$.

To summarize: the James-Stein estimator is essentially optimal over all linear estimators. Moreover, it is asymptotically optimal over all estimators, not just linear estimators. This also shows that the minimax risk and the linear minimax risk are asymptotically equivalent. This turns out to (sometimes) be a more general phenomenon, as we shall see.

### 7.7 Adaptive Estimation Over Sobolev Spaces

Theorem 7.32 gives an estimator that is minimax over $\Theta(m, c)$. However, the estimator is unsatisfactory because it requires that we know $c$ and $m$.

Efromovich and Pinsker (1984) proved the remarkable result that there exists an estimator that is minimax over $\Theta(m, c)$ without requiring knowledge of $m$ or $c$. The estimator is said to be adaptively asymptotically minimax. The idea is to divide the observations into blocks $B_{1}=\left\{Z_{1}, \ldots, Z_{n_{1}}\right\}, B_{2}=$ $\left\{Z_{n_{1}+1}, \ldots, Z_{n_{2}}\right\}, \ldots$ and then apply a suitable estimation procedure within blocks.

Here is particular block estimation scheme due to Cai et al. (2000). For any real number $a$ let $[a]$ denote the integer part of $a$. Let $b=1+1 / \log n$ and let $K_{0}$ be an integer such that $\left[b^{K_{0}}\right] \geq 3$ and $\left[b^{k}\right]-\left[b^{k-1}\right] \geq 3$ for $k \geq K_{0}+1$. Let $B_{0}=\left\{Z_{i}: 1 \leq i \leq\left[b^{K_{0}}\right]\right\}$ and let $B_{k}=\left\{Z_{i}: \quad\left[b^{k-1}\right]<i \leq\left[b^{k}\right]\right\}$ for $k \geq K_{0}+1$. Let $\widehat{\theta}$ be the estimator obtained by applying the James-Stein estimator within each block $B_{k}$. The estimator is taken to be 0 for $i>\left[b^{K_{1}}\right]$ where $K_{1}=\left[\log _{b}(n)\right]-1$.
7.50 Theorem (Cai, Low and Zhao, 2000). Let $\widehat{\theta}$ be the estimator above. Let $\Theta(m, c)=\left\{\theta: \quad \sum_{i=1}^{\infty} a_{i}^{2} \theta_{i}^{2} \leq c^{2}\right\}$ where $a_{1}=1$ and $a_{2 i}=a_{2 i+1}=1+(2 i \pi)^{2 m}$. Let $R_{n}(m, c)$ denote the minimax risk over $\Theta(m, c)$. Then for all $m>0$ and $c>0$,

$$
\lim _{n \rightarrow \infty} \frac{\sup _{\theta \in \Theta(m, c)} R(\widehat{\theta}, \theta)}{R_{n}(m, c)}=1
$$

### 7.8 Confidence Sets

In this section we discuss the construction of confidence sets for $\theta^{n}$. It will now be convenient to write $\theta$ and $Z$ instead of $\theta^{n}$ and $Z^{n}$.

Recall that $\mathcal{B}_{n} \subset \mathbb{R}^{n}$ is a $1-\alpha$ confidence set if

$$
\begin{equation*}
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) \geq 1-\alpha \tag{7.51}
\end{equation*}
$$

We have written the probability distribution $\mathbb{P}_{\theta}$ with the subscript $\theta$ to emphasize that the distribution depends on $\theta$. Here are some methods for constructing confidence sets.

Method I: The $\chi^{2}$ Confidence Set. The simplest confidence set for $\theta$ is based on the fact that $\|Z-\theta\|^{2} / \sigma_{n}^{2}$ has a $\chi_{n}^{2}$ distribution. Let

$$
\begin{equation*}
\mathcal{B}_{n}=\left\{\theta \in \mathbb{R}^{n}:\|Z-\theta\|^{2} \leq \sigma_{n}^{2} \chi_{n, \alpha}^{2}\right\} \tag{7.52}
\end{equation*}
$$

where $\chi_{n, \alpha}^{2}$ is the upper $\alpha$ quantile of a $\chi^{2}$ random variable with $n$ degrees of freedom. It follows immediately that

$$
\mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right)=1-\alpha, \quad \text { for all } \theta \in \mathbb{R}^{n} .
$$

Hence, $(7.51)$ is satisfied. The expected radius of this ball is $n \sigma_{n}^{2}$. We will see that we can improve on this.

Improving the $\chi^{2}$ Ball by Pre-testing. Before discussing more complicated methods, here is a simple idea-based on ideas in Lepski (1999)for improving the $\chi^{2}$ ball. The methods that follow are generalizations of this method.

Note that the $\chi^{2}$ ball $\mathcal{B}_{n}$ has a fixed radius $s_{n}=\sigma_{n} \sqrt{n}$. When applied to function estimation, $\sigma_{n}=O(1 / \sqrt{n})$ so that $s_{n}=O(1)$ and hence the radius of the ball does not even converge to zero as $n \rightarrow \infty$. The following construction makes the radius smaller. The idea is to test the hypothesis that $\theta=\theta_{0}$. If we accept the null hypothesis, we use a smaller ball centered at $\theta_{0}$. Here are the details.

First, test the hypothesis that $\theta=(0, \ldots, 0)$ using $\sum_{i} Z_{i}^{2}$ as a test statistic. Specifically, reject the null when

$$
T_{n}=\sum_{i} Z_{i}^{2}>c_{n}^{2}
$$

and $c_{n}$ is defined by

$$
\mathbb{P}\left(\chi_{n}^{2}>\frac{c_{n}^{2}}{\sigma_{n}^{2}}\right)=\frac{\alpha}{2}
$$

By construction, the test has type one error rate $\alpha / 2$. If $Z$ denotes a $\mathrm{N}(0,1)$ random variable then

$$
\frac{\alpha}{2}=\mathbb{P}\left(\chi_{n}^{2}>\frac{c_{n}^{2}}{\sigma_{n}^{2}}\right)=\mathbb{P}\left(\frac{\chi_{n}^{2}-n}{\sqrt{2 n}}>\frac{\frac{c_{n}^{2}}{\sigma_{n}^{2}}-n}{\sqrt{2 n}}\right) \approx \mathbb{P}\left(Z>\frac{\frac{c_{n}^{2}}{\sigma_{n}^{2}}-n}{\sqrt{2 n}}\right)
$$

implying that

$$
c_{n}^{2} \approx \sigma_{n}^{2}\left(n+\sqrt{2 n} z_{\alpha / 2}\right)
$$

Now we compute the power of this test when $\|\theta\|>\Delta_{n}$ where

$$
\Delta_{n}=\sqrt{2 \sqrt{2} z_{\alpha / 2}} n^{1 / 4} \sigma_{n}
$$

Write $Z_{i}=\theta_{i}+\sigma_{n} \epsilon_{i}$ where $\epsilon_{i} \sim N(0,1)$. Then,

$$
\begin{aligned}
\mathbb{P}_{\theta}\left(T_{n}>c_{n}^{2}\right) & =\mathbb{P}_{\theta}\left(\sum_{i} Z_{i}^{2}>c_{n}^{2}\right)=\mathbb{P}_{\theta}\left(\sum_{i}\left(\theta_{i}+\sigma_{n} \epsilon_{i}\right)^{2}>c_{n}^{2}\right) \\
& =\mathbb{P}_{\theta}\left(\|\theta\|^{2}+2 \sigma_{n} \sum_{i} \theta_{i} \epsilon_{i}+\sigma_{n}^{2} \sum_{i} \epsilon_{i}^{2}>c_{n}^{2}\right) .
\end{aligned}
$$

Now, $\|\theta\|^{2}+2 \sigma_{n} \sum_{i} \theta_{i} \epsilon_{i}+\sigma_{n}^{2} \sum_{i} \epsilon_{i}^{2}$ has mean $\|\theta\|^{2}+n \sigma_{n}^{2}$ and variance $4 \sigma_{n}^{2}\|\theta\|^{2}+$ $2 n \sigma_{n}^{4}$. Hence, with $Z$ denoting a $\mathrm{N}(0,1)$ random variable,

$$
\begin{aligned}
\mathbb{P}_{\theta}\left(T_{n}>c_{n}^{2}\right) & \approx \mathbb{P}\left(\|\theta\|^{2}+n \sigma_{n}^{2}+\sqrt{4 \sigma_{n}^{2}\|\theta\|^{2}+2 n \sigma_{n}^{4}} Z>c_{n}^{2}\right) \\
& \approx \mathbb{P}\left(\|\theta\|^{2} n \sigma_{n}^{2}+\sqrt{4 \sigma_{n}^{2}\|\theta\|^{2}+2 n \sigma_{n}^{4}} Z>\sigma_{n}^{2}\left(n+\sqrt{2 n} z_{\alpha / 2}\right)\right) \\
& =\mathbb{P}\left(Z>\frac{\left(\sqrt{2} z_{\alpha / 2}-\frac{\|\theta\|^{2}}{\sqrt{n} \sigma_{n}^{2}}\right)}{2+\frac{4\|\theta\|^{2}}{n \sigma_{n}^{2}}}\right) \geq \mathbb{P}\left(Z>\frac{\left(\sqrt{2} z_{\alpha / 2}-\frac{\|\theta\|^{2}}{\sqrt{n} \sigma_{n}^{2}}\right)}{2}\right) \\
& \geq 1-\frac{\alpha}{2}
\end{aligned}
$$

since $\|\theta\|>\Delta_{n}$ implies that

$$
\frac{\left(\sqrt{2} z_{\alpha / 2}-\frac{\|\theta\|^{2}}{\sqrt{n} \sigma_{n}^{2}}\right)}{2} \geq-z_{\alpha / 2}
$$

In summary, the test has type-one error $\alpha / 2$ and type-two error no more than $\alpha / 2$ for all $\|\theta\|>\Delta_{n}$.

Next we define the confidence procedure as follows. Let $\phi=0$ if the test accepts and $\phi=1$ if the test rejects. Define

$$
R_{n}=\left\{\begin{array}{cl}
\mathcal{B}_{n} & \text { if } \phi=1 \\
\left\{\theta:\|\theta\| \leq \Delta_{n}\right\} & \text { if } \phi=0 .
\end{array}\right.
$$

Thus, $R_{n}$ is a random radius confidence ball. The radius is the same as the $\chi^{2}$ ball when $\phi=1$ but when $\phi=0$, the radius is $\Delta_{n}$ which is much smaller. Let us now verify that the ball has the right coverage.

The noncoverage of this ball when $\theta=(0, \ldots, 0)$ is

$$
\begin{aligned}
\mathbb{P}_{0}(\theta \notin R) & =\mathbb{P}_{0}(\theta \notin R, \phi=0)+\mathbb{P}_{0}(\theta \notin R, \phi=1) \\
& \leq 0+\mathbb{P}_{0}(\phi=1)=\frac{\alpha}{2}
\end{aligned}
$$

The noncoverage of this ball when $\theta \neq(0, \ldots, 0)$ and $\|\theta\| \leq \Delta_{n}$ is

$$
\begin{aligned}
\mathbb{P}_{\theta}(\theta \notin R) & =\mathbb{P}_{\theta}(\theta \notin R, \phi=0)+\mathbb{P}_{\theta}(\theta \notin R, \phi=1) \\
& \leq 0+\mathbb{P}_{\theta}(\theta \notin B)=\frac{\alpha}{2}
\end{aligned}
$$

The noncoverage of this ball when $\theta \neq(0, \ldots, 0)$ and $\|\theta\|>\Delta_{n}$ is

$$
\begin{aligned}
\mathbb{P}_{\theta}(\theta \notin R) & =\mathbb{P}_{\theta}(\theta \notin R, \phi=0)+\mathbb{P}_{\theta}(\theta \notin R, \phi=1) \\
& \leq \mathbb{P}_{\theta}(\phi=0)+\mathbb{P}_{\theta}(\theta \notin B) \leq \frac{\alpha}{2}+\frac{\alpha}{2}=\alpha
\end{aligned}
$$

In summary, by testing whether $\theta$ is close to $(0, \ldots, 0)$ and using a smaller ball centered at $(0, \ldots, 0)$ when the test accepts, we get a ball with proper coverage and whose radius is sometimes smaller than the $\chi^{2}$ ball. The message is that:
a random radius confidence ball can have an expected radius that is smaller than a fixed confidence ball at some points in the parameter space.

The next section generalizes this idea.

Method II: The Baraud Confidence Set. Here we discuss the method due to Baraud (2004) which builds on Lepski (1999), as discussed above. We begin with a class $\mathcal{S}$ of linear subspaces of $\mathbb{R}^{n}$. Let $\Pi_{S}$ denote the projector onto $S$. Thus, for any vector $Z \in \mathbb{R}^{n}, \Pi_{S} Z$ is the vector in $S$ closest to $Z$.

For each subspace $S$, we construct a ball $\mathcal{B}_{S}$ of radius $\rho_{S}$ centered at an estimator in $S$, namely,

$$
\begin{equation*}
\mathcal{B}_{S}=\left\{\theta:\left\|\theta-\Pi_{S} Z\right\| \leq \rho_{S}\right\} \tag{7.53}
\end{equation*}
$$

For each $S \in \mathcal{S}$, we test whether $\theta$ is close to $S$ using $\left\|Z-\Pi_{S} Z\right\|$ as a test statistic. We then take the smallest confidence ball $\mathcal{B}_{S}$ among all unrejected
subspaces $S$. The key to making this work is this: the radius $\rho_{S}$ is chosen so that

$$
\begin{equation*}
\max _{\theta} \mathbb{P}_{\theta}\left(S \text { is not rejected and } \theta \notin \mathcal{B}_{S}\right) \leq \alpha_{S} \tag{7.54}
\end{equation*}
$$

where $\sum_{S \in \mathcal{S}} \alpha_{S} \leq \alpha$. The resulting confidence ball has coverage at least $1-\alpha$ since

$$
\begin{aligned}
\max _{\theta} \mathbb{P}_{\theta}(\theta \notin \mathcal{B}) & \leq \sum_{S} \max _{\theta} \mathbb{P}_{\theta}\left(S \text { is not rejected and } \theta \notin \mathcal{B}_{S}\right) \\
& =\sum_{S} \alpha_{S} \leq \alpha
\end{aligned}
$$

We will see that the $n$-dimensional maximization over $\theta \in \mathbb{R}^{n}$ can be reduced to a one-dimensional maximization since the probabilities only depend on $\theta$ through the quantity $z=\left\|\theta-\Pi_{S} \theta\right\|$.

The confidence set has coverage $1-\alpha$ even if $\theta$ is not close to one of the subspaces in $\mathcal{S}$. However, if it is close to one of the subspaces in $\mathcal{S}$, then the confidence ball will be smaller than the $\chi^{2}$ ball.

For example, suppose we expand a function $f(x)=\sum_{j} \theta_{j} \phi_{j}(x)$ in a basis, as in Section 7.3. Then, the $\theta_{i} \mathrm{~s}$ correspond to the coefficients of $f$ in this basis. If the function is smooth, then we expect that $\theta_{i}$ will be small for large $i$. Hence, $\theta$ might be well approximated by a vector of the form $\left(\theta_{1}, \ldots, \theta_{m}, 0, \ldots, 0\right)$. This suggests that we could test whether $\theta$ is close to the subspace $S_{m}$ of the vectors of the form $\left(\theta_{1}, \ldots, \theta_{m}, 0, \ldots, 0\right)$, for $m=0, \ldots, n$. In this case we would take the class of subspaces to be $\mathcal{S}=\left\{S_{0}, \ldots, S_{n}\right\}$.

Before we proceed with the details, we need some notation. If $X_{j} \sim N\left(\mu_{j}, 1\right)$, $j=1, \ldots, k$ are IID, then $T=\sum_{j=1}^{k} X_{j}^{2}$ has a noncentral $\chi^{2}$ distribution with noncentrality parameter $d=\sum_{j} \mu_{j}^{2}$ and $k$ degrees of freedom and we write $T \sim \chi_{d, k}^{2}$. Let $G_{d, k}$ denote the CDF of this random variable and let $q_{d, k}(\alpha)=G_{d, k}^{-1}(1-\alpha)$ denote the upper $\alpha$ quantile. By convention, we define $q_{d, k}(\alpha)=-\infty$ for $\alpha \geq 1$.

Let $\mathcal{S}$ be a finite collection of linear subspaces of $\mathbb{R}^{n}$. We assume that $\mathbb{R}^{n} \in \mathcal{S}$. Let $d(S)$ be the dimension of $S \in \mathcal{S}$ and let $e(S)=n-d(S)$. Fix $\alpha \in(0,1)$ and $\gamma \in(0,1)$ where $\gamma<1-\alpha$. Let

$$
\begin{equation*}
\mathcal{A}=\left\{S: \frac{\left\|Z-\Pi_{S} Z\right\|^{2}}{\sigma_{n}^{2}} \leq c(S)\right\} \tag{7.55}
\end{equation*}
$$

where

$$
\begin{equation*}
c(S)=q_{0, e(S)}(\gamma) \tag{7.56}
\end{equation*}
$$

Think of $\left\|Z-\Pi_{S} Z\right\|^{2}$ as a test statistic for the hypothesis that $\theta \in S$. Then $\mathcal{A}$ is the set of nonrejected subspaces. Note that $\mathcal{A}$ always includes the subspace $S=\mathbb{R}^{n}$ since, when $S=\mathbb{R}^{n}, \Pi_{S} Z=Z$ and $\left\|Z-\Pi_{S} Z\right\|^{2}=0$.

Let $\left(\alpha_{S}: S \in \mathcal{S}\right)$ be a set of numbers such that $\sum_{S \in \mathcal{S}} \alpha_{S} \leq \alpha$. Now define the $\rho_{S}$ as follows:

$$
\rho_{S}^{2}=\sigma_{n}^{2} \times \begin{cases}\inf _{z>0}\left\{G_{z, n}\left(q_{0, n}(\gamma)\right) \leq \alpha_{S}\right\} & \text { if } d(S)=0  \tag{7.57}\\ \sup _{z>0}\left\{z+q_{0, d(S)}\left(\frac{\alpha_{S}}{G_{z, e(S)}(c(S))}\right)\right\} & \text { if } 0<d(S)<n \\ \rho_{S}^{2}=\sigma_{n}^{2} q_{0, n}\left(\alpha_{S}\right) & \text { if } d(S)=n\end{cases}
$$

Define

$$
\widehat{S}=\operatorname{argmin}_{S \in \mathcal{A}} \rho_{S}
$$

$\widehat{\theta}=\Pi_{\widehat{S}} Z$, and $\widehat{\rho}=\rho_{\widehat{S}}$. Finally, define

$$
\begin{equation*}
\mathcal{B}_{n}=\left\{\theta \in \mathbb{R}^{n}:\|\theta-\widehat{\theta}\|^{2} \leq \widehat{\rho}^{2}\right\} \tag{7.58}
\end{equation*}
$$

7.59 Theorem (Baraud 2004). The set $\mathcal{B}_{n}$ defined in (7.58) is a valid confidence set:

$$
\begin{equation*}
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) \geq 1-\alpha \tag{7.60}
\end{equation*}
$$

Proof. Let $\mathcal{B}_{S}=\left\{\theta:\left\|\theta-\Pi_{S} Z\right\|^{2} \leq \rho_{S}^{2}\right\}$. Then,

$$
\begin{aligned}
\mathbb{P}_{\theta}\left(\theta \notin \mathcal{B}_{n}\right) & \leq \mathbb{P}_{\theta}\left(\theta \notin \mathcal{B}_{S} \text { for some } S \in \mathcal{A}\right) \\
& \leq \sum_{S} \mathbb{P}_{\theta}\left(\left\|\theta-\Pi_{S} Z\right\|>\rho_{S}, \widehat{S} \in \mathcal{A}\right) \\
& =\sum_{S} \mathbb{P}_{\theta}\left(\left\|\theta-\Pi_{S} Z\right\|>\rho_{S},\left\|Z-\Pi_{S} Z\right\|^{2} \leq c(S) \sigma_{n}^{2}\right)
\end{aligned}
$$

Since $\sum_{S} \alpha_{S} \leq \alpha$, it suffices to show that $a(S) \leq \alpha_{S}$ for all $S \in \mathcal{S}$, where

$$
\begin{equation*}
a(S) \equiv \mathbb{P}_{\theta}\left(\left\|\theta-\Pi_{S} Z\right\|>\rho_{S},\left\|Z-\Pi_{S} Z\right\|^{2} \leq \sigma_{n}^{2} c(S)\right) \tag{7.61}
\end{equation*}
$$

When $d(S)=0, \Pi_{S} Z=(0, \ldots, 0)$. If $\|\theta\| \leq \rho_{S}$ then $a(0)=0$ which is less than $\alpha_{S}$. If $\|\theta\|>\rho_{S}$, then

$$
\begin{aligned}
a(S) & =\mathbb{P}_{\theta}\left(\sum_{i=1}^{n} Z_{i}^{2} \leq \sigma_{n}^{2} q_{0, n}(\gamma)\right) \\
& =G_{\|\theta\|^{2} / \sigma_{n}^{2}, n}\left(q_{0, n}(\gamma)\right) \leq G_{\rho_{0}^{2} / \sigma_{n}^{2}, n}\left(q_{0, n}(\gamma)\right) \\
& \leq \alpha_{S}
\end{aligned}
$$

since $G_{d, n}(u)$ is decreasing in $z$ for all $u$ and from the definition of $\rho_{0}^{2}$.
Now consider the case where $0<d(S)<n$. Let

$$
A=\frac{\left\|\theta-\Pi_{S} Z\right\|^{2}}{\sigma_{n}^{2}}=z+\sum_{j=1}^{m} \epsilon_{j}^{2}, \quad B=\frac{\|\widehat{\theta}-Z\|^{2}}{\sigma_{n}^{2}}
$$

where $z=\left\|\theta-\Pi_{S} \theta\right\|^{2} / \sigma_{n}^{2}$. Then $A$ and $B$ are independent, $A \sim z+\chi_{0, d(S)}^{2}$, and $B \sim \chi_{z, e(S)}^{2}$. Hence,

$$
\begin{align*}
a(S) & =\mathbb{P}_{\theta}\left(A>\frac{\rho_{m}^{2}}{\sigma_{n}^{2}}, B<c(S)\right) \\
& =\mathbb{P}_{\theta}\left(z+\chi_{d(S)}^{2}>\frac{\rho_{S}^{2}}{\sigma_{n}^{2}}, \chi_{z, e(S)}^{2}<c(S)\right)  \tag{7.62}\\
& =\left(1-G_{0, d(S)}\left(\frac{\rho_{S}^{2}}{\sigma_{n}^{2}}-z\right)\right) \times G_{z, e(S)}(c(S)) . \tag{7.63}
\end{align*}
$$

From the definition of $\rho_{S}^{2}$,

$$
\frac{\rho_{S}^{2}}{\sigma_{n}^{2}}-z \geq q_{0, d(S)}\left(\frac{\alpha_{S}}{G_{z, e(S)}(c(S))} \wedge 1\right)
$$

and hence,

$$
\begin{align*}
1-G_{0, d(S)}\left(\frac{\rho_{S}^{2}}{\sigma_{n}^{2}}-z\right) & \leq 1-G_{0, d(S)}\left(q_{0, d(S)}\left(\frac{\alpha_{S}}{G_{z, e(S)}(c(S))}\right)\right) \\
& =\frac{\alpha_{S}}{G_{z, e(S)}(c(S))} \tag{7.64}
\end{align*}
$$

It then follows (7.63) and (7.64) that $a(S) \leq \alpha_{S}$.
For the case $d(S)=n, \Pi_{S} Z=Z$, and $\left\|\theta-\Pi_{S} Z\right\|^{2}=\sigma_{n}^{2} \sum_{i=1}^{n} \epsilon_{i}^{2} \stackrel{d}{=} \sigma_{n}^{2} \chi_{n}^{2}$ and so

$$
a(S)=\mathbb{P}_{\theta}\left(\sigma_{n}^{2} \chi_{n}^{2}>q_{0, n}\left(\alpha_{S}\right) \sigma_{n}^{2}\right)=\alpha_{S}
$$

by the definition of $q_{0, n}$.
When $\sigma_{n}$ is unknown we estimate the variance using one of the methods discussed in Chapter 5 and generally the coverage is only asymptotically correct. To see the effect of having uncertainty about $\sigma_{n}$, consider the idealized case where $\sigma_{n}$ is known to lie with certainty in the interval $I=\left[\sqrt{1-\eta_{n}} \tau_{n}, \tau_{n}\right]$. (In practice, we would construct a confidence interval for $\sigma$ and adjust the
level $\alpha$ of the confidence ball appropriately.) In this case, the radii $\rho_{S}$ are now defined by:

$$
\rho_{S}^{2}= \begin{cases}\inf _{z>0}\left\{\sup _{\sigma_{n} \in I} G_{z / \sigma_{n}^{2}, n}\left(q_{0, n}(\gamma) \tau_{n}^{2} / \sigma_{n}^{2}\right) \leq \alpha_{S}\right\} & \text { if } d(S)=0  \tag{7.65}\\ \sup _{z>0, \sigma_{n} \in I}\left\{z \sigma_{n}^{2}+\sigma_{n}^{2} q_{0, d(S)}\left(h_{S}\left(z, \sigma_{n}\right)\right)\right\} & \text { if } 0<d(S)<n \\ q_{0, n}\left(\alpha_{S}\right) \tau_{n}^{2} & \text { if } d(S)=n\end{cases}
$$

where

$$
\begin{equation*}
h_{S}(z, \sigma)=\frac{\alpha_{S}}{G_{z, e(S)}\left(G_{z, e(S)}\left(q_{0, e(S)}(\gamma) \tau_{n}^{2} / \sigma^{2}\right)\right)} \tag{7.66}
\end{equation*}
$$

and $\mathcal{A}$ is now defined by

$$
\begin{equation*}
\mathcal{A}=\left\{S \in \mathcal{S}:\left\|Z-\Pi_{S} Z\right\|^{2} \leq q_{0, e(S)}(\gamma) \tau_{n}^{2}\right\} \tag{7.67}
\end{equation*}
$$

Beran-Dümbgen-Stein Pivotal method. Now we discuss a different approach due to Stein (1981) and developed further by Li (1989), Beran and Dümbgen (1998), and Genovese and Wasserman (2005). The method is simpler than the Baraud-Lepski approach but it uses asymptotic approximations. This method is considered in more detail in the next chapter but here is the basic idea.

Consider nested subsets $\mathcal{S}=\left\{S_{0}, S_{1}, \ldots, S_{n}\right\}$ where

$$
S_{j}=\left\{\theta=\left(\theta_{1}, \ldots, \theta_{j}, 0, \ldots, 0\right):\left(\theta_{1}, \ldots, \theta_{j}\right) \in \mathbb{R}^{j}\right\}
$$

Let $\widehat{\theta}_{m}=\left(Z_{1}, \ldots, Z_{m}, 0, \ldots, 0\right)$ denote the estimator under model $S_{m}$. The loss function is

$$
L_{m}=\left\|\widehat{\theta}_{m}-\theta\right\|^{2}
$$

Define the pivot

$$
\begin{equation*}
V_{m}=\sqrt{n}\left(L_{m}-\widehat{R}_{m}\right) \tag{7.68}
\end{equation*}
$$

where $\widehat{R}_{m}=m \sigma_{n}^{2}+\sum_{j=m+1}^{n}\left(Z_{j}^{2}-\sigma_{n}^{2}\right)$ is SURE. Let $\widehat{m}$ minimize $\widehat{R}_{m}$ over $m$. Beran and Dümbgen (1998) show that $V_{\widehat{m}} / \widehat{\tau} \rightsquigarrow N(0,1)$ where

$$
\tau_{m}^{2}=\mathbb{V}\left(V_{m}\right)=2 n \sigma_{n}^{2}\left(n \sigma_{n}^{2}+2 \sum_{j=m+1}^{n} \theta_{j}^{2}\right)
$$

and

$$
\widehat{\tau}^{2}=2 n \sigma_{n}^{2}\left(n \sigma_{n}^{2}+2 \sum_{j=\widehat{m}+1}^{n}\left(Z_{j}^{2}-\sigma_{n}^{2}\right)\right)
$$

Let

$$
r_{n}^{2}=\widehat{R}_{m}+\frac{\widehat{\tau} z_{\alpha}}{\sqrt{n}}
$$

and define

$$
\mathcal{B}_{n}=\left\{\theta \in \mathbb{R}^{n}:\left\|\theta_{m}-\widehat{\theta}\right\|^{2} \leq r_{n}^{2}\right\} .
$$

Then,

$$
\begin{aligned}
\mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) & =\mathbb{P}_{\theta}\left(\|\theta-\widehat{\theta}\|^{2} \leq r_{n}^{2}\right)=\mathbb{P}_{\theta}\left(L_{m} \leq r_{n}^{2}\right) \\
& =\mathbb{P}_{\theta}\left(L_{m} \leq \widehat{R}_{m}+\frac{\widehat{\tau} z_{\alpha}}{\sqrt{n}}\right)=\mathbb{P}_{\theta}\left(\frac{V_{\widehat{m}}}{\widehat{\tau}} \leq z_{\alpha}\right) \\
& \rightarrow 1-\alpha
\end{aligned}
$$

A practical problem with this method is that $r_{n}^{2}$ can be negative. This is due to the presence of the term $\sum_{j=m+1}^{n}\left(Z_{j}^{2}-\sigma_{n}^{2}\right)$ in $\widehat{R}$ and $\tau$. We deal with this by replacing such terms with $\max \left\{\sum_{j=m+1}^{n}\left(Z_{j}^{2}-\sigma_{n}^{2}\right), 0\right\}$. This can lead to over-coverage but at least leads to well-defined radii.
7.69 Example. Consider nested subsets $\mathcal{S}=\left\{S_{0}, S_{1}, \ldots, S_{n}\right\}$ where $S_{0}=$ $\{(0, \ldots, 0)\}$ and

$$
S_{j}=\left\{\theta=\left(\theta_{1}, \ldots, \theta_{j}, 0, \ldots, 0\right):\left(\theta_{1}, \ldots, \theta_{j}\right) \in \mathbb{R}^{j}\right\}
$$

We take $\alpha=0.05, n=100, \sigma_{n}=1 / \sqrt{n}$, and $\alpha_{S}=\alpha /(n+1)$ for all $S$ so that $\sum \alpha_{S}=\alpha$ as required. Figure 7.2 shows $\rho_{S}$ versus the dimension of $S$ for $\gamma=0.05,0.15,0.50,0.90$. The dotted line is the radius of the $\chi^{2}$ ball. One can show that

$$
\begin{equation*}
\frac{\rho_{0}}{\rho_{n}}=O\left(n^{-1 / 4}\right) \tag{7.70}
\end{equation*}
$$

which shows that shrinking towards lower-dimensional models leads to smaller confidence sets. There is an interesting tradeoff. Setting $\gamma$ large makes $\rho_{0}$ small leading to a potentially smaller confidence ball. However, making $\gamma$ large increases the set $\mathcal{A}$ which diminishes the chances of choosing a small $\rho$. We simulated under the model $\theta=(10,10,10,10,10,0, \ldots, 0)$. See Table 7.1 for a summary. In this example, the pivotal method seems to perform the best.

### 7.9 Optimality of Confidence Sets

How small can we make the confidence set while still maintaining correct coverage? In this section we will see that if $\mathcal{B}_{n}$ is a confidence ball with radius


FIGURE 7.2. Constants $\rho_{S}$ from Example 7.69. The horizontal axis is the dimension of the submodel. The four curves show $\rho_{S}$ for $\gamma=0.05,0.15,0.50,0.90$. The highest curve corresponds to $\gamma=0.05$ and the curves get lower as $\gamma$ increases. The dotted line is the radius of the $\chi^{2}$ ball.

| Method | Coverage | Radius |
| :--- | :---: | :---: |
| $\chi^{2}$ | 0.950 | 1.115 |
| Baraud $(\gamma=0.90)$ | 1.000 | 0.973 |
| $(\gamma=0.50)$ | 1.000 | 0.904 |
| $(\gamma=0.15)$ | 1.000 | 0.779 |
| $(\gamma=0.05)$ | 0.996 | 0.605 |
| Pivotal | 0.998 | 0.582 |

TABLE 7.1. Simulation results from Example 7.69 based on 1000 simulations.
$s_{n}$ then $E_{\theta}\left(s_{n}\right) \geq C_{1} \sigma_{n} n^{1 / 4}$ for every $\theta$ and $E_{\theta}\left(s_{n}\right) \geq C_{2} \sigma_{n} n^{1 / 2}$ for some $\theta$. Here, $C_{1}$ and $C_{2}$ are positive constants. The $\chi^{2}$ ball has radius $\sigma_{n} n^{1 / 2}$ for all $\theta$. This suggests that the $\chi^{2}$ ball can be improved upon, and indeed, the Baraud confidence ball can achieve the faster $\sigma_{n} n^{1 / 4}$ rate at some points in the parameter space. We will provide some of the details in this section. But first, let us compare this with point estimation.

From Theorem 7.32, the optimal rate of convergence of a point estimator over a Sobolev space of order $m$ is $n^{-2 m /(2 m+1)}$. According to Theorem 7.50, we can construct estimators that achieve this rate, without prior knowledge of $m$. This raises the following questions: Can we construct confidence balls that adaptively achieve this optimal rate? The short answer is no. Robins and van der Vaart (2005), Juditsky and Lambert-Lacroix (2003), and Cai and Low (2005) show that some degree of adaptivity is possible for confidence sets but the amount of adaptivity is quite restricted. Without any smoothness assumptions, we see from our comments above that the fastest rate of convergence one can attain is $\sigma_{n} n^{1 / 4}$ which is of order $O\left(n^{-1 / 4}\right)$ when $\sigma_{n}=\sigma / \sqrt{n}$.

Turning to the details, we begin with the following Theorem due to Li (1989).
7.71 Theorem (Li 1989). Let $\mathcal{B}_{n}=\left\{\theta^{n} \in \mathbb{R}^{n}:\left\|\widehat{\theta}^{n}-\theta^{n}\right\| \leq s_{n}\right\}$ where $\widehat{\theta}^{n}$ is any estimator of $\theta^{n}$ and $s_{n}=s_{n}\left(Z^{n}\right)$ is the radius of the ball. Suppose that

$$
\begin{equation*}
\liminf _{n \rightarrow \infty} \inf _{\theta^{n} \in \mathbb{R}^{n}} \mathbb{P}_{\theta^{n}}\left(\theta^{n} \in \mathcal{B}_{n}\right) \geq 1-\alpha \tag{7.72}
\end{equation*}
$$

Then for any sequence $\theta^{n}$ and any $c_{n} \rightarrow 0$,

$$
\begin{equation*}
\limsup _{n \rightarrow \infty} \mathbb{P}_{\theta^{n}}\left(s_{n} \leq c_{n} \sigma_{n} n^{1 / 4}\right) \leq \alpha \tag{7.73}
\end{equation*}
$$

Finite sample results are available from Baraud (2004) and Cai and Low (2005). For example, we have the following result, whose proof is in the appendix.
7.74 Theorem (Cai and Low 2004). Assume the model (7.1). Fix $0<\alpha<$ $1 / 2$. Let $\mathcal{B}_{n}=\left\{\theta:\|\widehat{\theta}-\theta\| \leq s_{n}\right\}$ be such that

$$
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) \geq 1-\alpha
$$

Then, for every $0<\epsilon<(1 / 2)-\alpha$,

$$
\begin{equation*}
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{E}_{\theta}\left(s_{n}\right) \geq \sigma_{n}(1-2 \alpha-2 \epsilon) n^{1 / 4}\left(\log \left(1+\epsilon^{2}\right)\right)^{1 / 4} \tag{7.75}
\end{equation*}
$$

In particular, if $\sigma_{n}=\sigma / \sqrt{n}$, then

$$
\begin{equation*}
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{E}_{\theta}\left(s_{n}\right) \geq \frac{C}{n^{1 / 4}} \tag{7.76}
\end{equation*}
$$

where $C=\sigma(1-2 \alpha-2 \epsilon)\left(\log \left(1+\epsilon^{2}\right)\right)^{1 / 4}$.
The lower bound in the above theorem cannot be attained everywhere, as the next result shows.
7.77 Theorem (Cai and Low 2004). Assume the model (7.1). Fix $0<\alpha<$ $1 / 2$. Let $\mathcal{B}_{n}=\left\{\theta:\|\widehat{\theta}-\theta\| \leq s_{n}\right\}$ be such that

$$
\inf _{\theta \in \mathbb{R}^{n}} \mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) \geq 1-\alpha
$$

Then, for every $0<\epsilon<(1 / 2)-\alpha$,

$$
\begin{equation*}
\sup _{\theta \in \mathbb{R}^{n}} \mathbb{E}_{\theta}\left(s_{n}\right) \geq \epsilon \sigma_{n} z_{\alpha+2 \epsilon} \sqrt{n} \sqrt{\frac{\epsilon}{1-\alpha-\epsilon}} . \tag{7.78}
\end{equation*}
$$

In particular, if $\sigma_{n}=\sigma / \sqrt{n}$, then

$$
\begin{equation*}
\sup _{\theta \in \mathbb{R}^{n}} \mathbb{E}_{\theta}\left(s_{n}\right) \geq C \tag{7.79}
\end{equation*}
$$

where $C=\epsilon z_{\alpha+2 \epsilon} \sqrt{\epsilon /(1-\alpha-\epsilon)}$.
Despite these pessimistic sounding results, there is some potential for adaptation since the infimum in Theorem 7.74 is smaller than the supremum in Theorem 7.77. For example, the $\chi^{2}$ ball has radius $O\left(\sigma_{n} \sqrt{n}\right)$ but the lower bound in the above theorem is $O\left(\sigma_{n} n^{1 / 4}\right)$ suggesting that we can do better than the $\chi^{2}$ ball. This was the motivation for the Baraud and pivotal confidence sets. The Baraud confidence set does have a certain type of adaptivity: if $\theta \in S$ then $\widehat{\rho} \leq \rho_{S}$ with high probability. This follows easily from the way that the ball is defined. Let us formalize this as a lemma.
7.80 Lemma. Define $\mathcal{S}, \alpha, \gamma$ and $\left(\rho_{S}: S \in \mathcal{S}\right)$ as in Theorem 7.59. For each $S \in \mathcal{S}$,

$$
\begin{equation*}
\inf _{\theta \in S} \mathbb{P}_{\theta}\left(\widehat{\rho} \leq \rho_{S}\right) \geq 1-\gamma \tag{7.81}
\end{equation*}
$$

Baraud also gives the following results which show that his construction is essentially optimal. The first result gives a lower bound on any adaptive confidence ball. The result after that shows that the radius $\rho_{S}$ of his confidence set essentially achieves this lower bound.
7.82 Theorem (Baraud 2004). Suppose that $\widehat{\theta}=\widehat{\theta}(Z)$ and $r=r(Z)$ are such that $\mathcal{B}=\left\{\theta:\|\theta-\widehat{\theta}\|^{2} \leq r^{2}\right\}$ is a $1-\alpha$ confidence ball. Also suppose that $2 \alpha+\gamma<1-e^{-1 / 36}$ and that $d(S) \leq n / 2$. If

$$
\begin{equation*}
\inf _{\theta \in S} \mathbb{P}_{\theta}\left(r \leq r_{S}\right) \geq 1-\gamma \tag{7.83}
\end{equation*}
$$

then, for some $C=C(\alpha, \gamma)>0$,

$$
\begin{equation*}
r_{S}^{2} \geq C \sigma_{n}^{2} \max \{d(S), \sqrt{n}\} \tag{7.84}
\end{equation*}
$$

Taking $S$ to consist of a single point yields the same result as Theorem 7.74 and taking $S=\mathbb{R}^{n}$ yields the same result as Theorem 7.77.
7.85 Theorem (Baraud 2004). Define $\mathcal{S}, \alpha, \gamma$ and $\left(\rho_{S}: S \in \mathcal{S}\right)$ as in Theorem 7.59. Assume that $d(S) \leq n / 2$ for every $S \in \mathcal{S}$ except for $S=\mathbb{R}^{n}$. There exists a universal constant $C>0$ such that

$$
\begin{equation*}
\rho_{S}^{2} \leq C \sigma_{n}^{2} \max \left\{d(S), \sqrt{n \log \left(1 / \alpha_{S}\right)}, \log \left(1 / \alpha_{S}\right)\right\} \tag{7.86}
\end{equation*}
$$

When $\sigma_{n}$ is only known to lie in an interval $I=\left[\sqrt{1-\eta_{n}} \tau_{n}, \tau_{n}\right]$, Baraud shows that the lower bound (7.84) becomes

$$
\begin{equation*}
r_{S}^{2} \geq C \tau_{n}^{2} \max \left\{\eta_{n} n / 2, d(S)\left(1-\eta_{n}\right), \sqrt{n-d(S)}\left(1-\eta_{n}\right)\right\} \tag{7.87}
\end{equation*}
$$

which shows that information about $\sigma$ is crucial. Indeed, the best we realistically could hope for is to know $\sigma^{2}$ up to order $\eta_{n}=O\left(n^{-1 / 2}\right)$ in which case the lower bound is of order $\max \{\sqrt{n}, d(S)\}$.

### 7.10 Random Radius Bands?

We have seen that random radius confidence balls can be adaptive in the sense that they can be smaller than fixed radius confidence balls at some points in the parameter space. Is the same true for confidence bands? The answer is no, as follows from results in Low (1997). Actually, Low considers estimating a density $f$ at a single point $x$ but essentially the same results apply to regression and to confidence bands. He shows that any random radius confidence interval for $f(x)$ must have expected width at least as large as a fixed width confidence interval. Thus, there is a qualitative difference between constructing a confidence ball versus a confidence band.

Similar comments apply for other norms. The $L_{p}$ norm is defined by

$$
\|\theta\|_{p}= \begin{cases}\left(\sum_{i}\left|\theta_{i}\right|^{p}\right)^{1 / p} & p<\infty \\ \max _{i}\left|\theta_{i}\right| & p=\infty\end{cases}
$$

Confidence bands can be thought of as $L_{\infty}$ confidence balls. It can be shown that confidence balls in the $L_{p}$ norm with $2<p<\infty$ fall in between the two extremes of $L_{2}$ and $L_{\infty}$ in the sense that they have some adaptivity, but not as much as in the $L_{2}$ norm. Similar comments apply to hypothesis testing; see Ingster and Suslina (2003).

### 7.11 Penalization, Oracles and Sparsity

Consider again the many Normal means problem

$$
Z_{i} \sim \theta_{i}+\sigma_{n} \epsilon_{i}, \quad i=1, \ldots, n
$$

If we choose $\widehat{\theta}$ to minimize the sums of squares $\sum_{i=1}^{n}\left(Z_{i}-\widehat{\theta}_{i}\right)^{2}$, we get the MLE $\widehat{\theta}=Z=\left(Z_{1}, \ldots, Z_{n}\right)$. If instead we minimize a penalized sums of squares, we get different estimators.
7.88 Theorem. Let $J: \mathbb{R}^{n} \rightarrow[0, \infty), \lambda \geq 0$ and define the penalized sums of squares

$$
M=\sum_{i=1}^{n}\left(Z_{i}-\theta_{i}\right)^{2}+\lambda J(\theta)
$$

Let $\widehat{\theta}$ minimize $M$. If $\lambda=0$ then $\widehat{\theta}=Z$. If $J(\theta)=\sum_{i=1}^{n} \theta_{i}^{2}$ then $\widehat{\theta}_{i}=Z_{i} /(1+\lambda)$ which is a linear shrinkage estimator. If $J(\theta)=\sum_{i=1}^{n}\left|\theta_{i}\right|$ then $\widehat{\theta}$ is the softthresholding estimator (7.21). If $J(\theta)=\#\left\{\theta_{i}: \quad \theta_{i} \neq 0\right\}$ then $\widehat{\theta}$ is the hardthresholding estimator (7.23).

Thus we see that linear shrinkage, soft thresholding and hard thresholding are all special cases of one general approach. The case of the $L_{1}$ penalty $\sum_{i=1}^{n}\left|\theta_{i}\right|$ is especially interesting. According to Theorem 7.88, the estimator that minimizes

$$
\begin{equation*}
\sum_{i=1}^{n}\left(Z_{i}-\widehat{\theta}_{i}\right)^{2}+\lambda \sum_{i=1}^{n}\left|\theta_{i}\right| \tag{7.89}
\end{equation*}
$$

is the soft-threshold estimator $\widehat{\theta}_{\lambda}=\left(\widehat{\theta}_{\lambda, 1}, \ldots, \widehat{\theta}_{\lambda, n}\right)$ where

$$
\widehat{\theta}_{i, \lambda}=\operatorname{sign}\left(Z_{i}\right)\left(\left|Z_{i}\right|-\lambda\right)_{+} .
$$

The criterion (7.89) arises in variable selection for linear regression under the name lasso (Tibshirani (1996)) and in signal processing under the name basis pursuit (Chen et al. (1998)). We will see in Chapter 9 that soft thresholding also plays an important role in wavelet methods.

To get more insight on soft thresholding, we consider a result from Donoho and Johnstone (1994). Consider estimating $\theta_{i}$ and suppose we use either $Z_{i}$ or 0 as an estimator. Such an estimator might be appropriate if we think the vector $\theta$ is sparse in the sense that it has many zeroes. The risk of $Z_{i}$ is $\sigma_{n}^{2}$ and the risk of 0 is $\theta_{i}^{2}$. Imagine an oracle that knows when $Z_{i}$ has better risk and when 0 has better risk. The risk of the oracles estimator is $\min \left\{\sigma_{n}^{2}, \theta_{i}^{2}\right\}$. The risk for estimating the whole vector $\theta$ is

$$
R_{\mathrm{oracle}}=\sum_{i=1}^{n} \min \left\{\sigma_{n}^{2}, \theta_{i}^{2}\right\}
$$

Donoho and Johnstone (1994) showed that soft thresholding gives an estimator that comes close to the oracle.
7.90 Theorem (Donoho and Johnstone 1994). Let $\lambda=\sigma_{n} \sqrt{2 \log n}$. Then, for every $\theta \in \mathbb{R}^{n}$,

$$
\mathbb{E}_{\theta}\left\|\widehat{\theta}_{\lambda}-\theta\right\|^{2} \leq(2 \log n+1)\left(\sigma_{n}^{2}+R_{\text {oracle }}\right)
$$

Moreover, no estimator can get substantially closer to the oracle in the sense that, as $n \rightarrow \infty$,

$$
\begin{equation*}
\inf _{\widehat{\theta}} \sup _{\theta \in \mathbb{R}^{n}} \frac{\mathbb{E}_{\theta}\|\widehat{\theta}-\theta\|^{2}}{\sigma_{n}^{2}+R_{\text {oracle }}} \sim 2 \log n . \tag{7.91}
\end{equation*}
$$

Consider now a sparse vector $\theta$ that is 0 except for $k$ large components, where $k \ll n$. Then, $R_{\text {oracle }}=k \sigma_{n}^{2}$. In function estimation problems, we will see in the next chapter that $\sigma_{n}^{2}=O(1 / n)$ and hence $R_{\text {oracle }}=O(k / n)$ which is small in sparse cases ( $k$ small).

### 7.12 Bibliographic Remarks

The idea of reducing nonparametric models to Normal means models (or the white noise model in the appendix) dates back at least to Ibragimov and Has'minskii (1977), Efromovich and Pinsker (1982), and others. See Brown and Low (1996), Nussbaum (1996a) for examples of recent results in this area. A thorough treatment of Normal decision theory and its relation to nonparametric problems is contained in Johnstone (2003). There is also a substantial literature on hypothesis testing in this framework. Many of the results are due to Ingster and are summarized in Ingster and Suslina (2003).

### 7.13 Appendix

The White Noise Model. Regression is also connected with the white noise model. Here is a brief description. Recall that a standard Brownian motion $W(t), 0 \leq t \leq 1$ is a random function such that $W(0)=0, W(s+t)-W(s) \sim$ $N(0, t)$ and, $W(v)-W(u)$ is independent of $W(t)-W(s)$ for $0 \leq u \leq v \leq$ $s \leq t$. You can think of $W$ as a continuous version of a random walk. Let $Z_{i}=f(i / n)+\sigma \epsilon_{i}$ with $\epsilon_{i} \sim N(0,1)$. For $0 \leq t \leq 1$, define

$$
Z_{n}(t)=\frac{1}{n} \sum_{i=1}^{[n t]} Z_{i}=\frac{1}{n} \sum_{i=1}^{[n t]} f(i / n)+\frac{\sigma}{\sqrt{n}} \frac{1}{\sqrt{n}} \sum_{i=1}^{[n t]} Z_{i} .
$$

The term $\frac{1}{n} \sum_{i=1}^{[n t]} f(i / n)$ converges to $\int_{0}^{t} f(s) d s$ as $n \rightarrow \infty$. The term $n^{-1 / 2}$ $\sum_{i=1}^{[n t]} Z_{i}$ converges to a standard Brownian motion. (For any fixed $t$, this is just an application of the central limit theorem.) Thus, asymptotically we can write

$$
Z(t)=\int_{0}^{t} f(s) d s+\frac{\sigma}{\sqrt{n}} W(t)
$$

This is called the standard white noise model, often written in differential form as

$$
\begin{equation*}
d Z(t)=f(t) d t+\frac{\sigma}{\sqrt{n}} d W(t) \tag{7.92}
\end{equation*}
$$

where $d W(t)$ is the white noise process. ${ }^{6}$
Let $\phi_{1}, \phi_{2}, \ldots$ be an orthonormal basis for $L_{2}(0,1)$ and write $f(x)=$ $\sum_{i=1}^{\infty} \theta_{i} \phi_{i}(x)$ where $\theta_{i}=\int f(x) \phi_{i}(x) d x$. Multiply (7.92) by $\phi_{j}$ and integrate. This yields $Z_{i}=\theta_{i}+(\sigma / \sqrt{n}) \epsilon_{i}$ where $Z_{i}=\int \phi_{i}(t) d Z(t)$ and $\epsilon_{i}=$ $\int \phi_{i}(t) d W(t) \sim N(0,1)$. We are back to the Normal means problem. A more complicated argument can be used to relate density estimation to the white noise model as in Nussbaum (1996a).

Weak Differentiability. Let $f$ be integrable on every bounded interval. Then $f$ is weakly differentiable if there exists a function $f^{\prime}$ that is integrable on every bounded interval, such that

$$
\int_{x}^{y} f^{\prime}(s) d s=f(y)-f(x)
$$

whenever $x \leq y$. We call $f^{\prime}$ the weak derivative of $f$. An equivalent condition is that for every $\phi$ that is compactly supported and infinitely differentiable,

$$
\int f(s) \phi^{\prime}(s) d s=-\int f^{\prime}(s) \phi(s) d s
$$

[^17]See Härdle et al. (1998), page 72.
Proof of Pinsker's Theorem (Theorem 7.28). (Following Nussbaum (1996b).) We will need to use Bayes estimators, which we now review. Let $\pi_{n}$ be a prior for $\theta^{n}$. The integrated risk is defined to be $B\left(\widehat{\theta}, \pi_{n}\right)=\int R\left(\widehat{\theta}^{n}, \theta^{n}\right) d \pi_{n}\left(\theta^{n}\right)=$ $\mathbb{E}_{\pi_{n}} \mathbb{E}_{\theta} L(\widehat{\theta}, \theta)$. The Bayes estimator $\widehat{\theta}_{\pi_{n}}$ minimizes the Bayes risk:

$$
\begin{equation*}
B\left(\pi_{n}\right)=\inf _{\hat{\theta}} B\left(\widehat{\theta}^{n}, \pi_{n}\right) . \tag{7.93}
\end{equation*}
$$

An explicit formula for the Bayes estimator is

$$
\widehat{\theta}_{\pi_{n}}(y)=\operatorname{argmin}_{a} \mathbb{E}\left(L(a, \theta) \mid Z^{n}\right) .
$$

In the case of squared error loss $L(a, \theta)=\|a-\theta\|_{n}^{2}$, the Bayes estimator is $\widehat{\theta}_{\pi_{n}}(y)=\mathbb{E}\left(\theta \mid Z^{n}\right)$.

Let $\Theta_{n}=\Theta_{n}(c)$. Let

$$
R_{n}=\inf _{\widehat{\theta}} \sup _{\theta \in \Theta_{n}} R(\widehat{\theta}, \theta)
$$

denote the minimax risk. We will find an upper bound and a lower bound on the risk.

Upper Bound. Let $\widehat{\theta}_{j}=c^{2} Z_{j} /\left(\sigma^{2}+c^{2}\right)$. The bias of this estimator is

$$
\mathbb{E}_{\theta}\left(\widehat{\theta}_{j}\right)-\theta_{j}=-\frac{\sigma^{2} \theta_{j}}{\sigma^{2}+c^{2}}
$$

and the variance is

$$
\mathbb{V}_{\theta}\left(\widehat{\theta}_{j}\right)=\left(\frac{c^{2}}{c^{2}+\sigma^{2}}\right)^{2} \sigma_{n}^{2}=\left(\frac{c^{2}}{c^{2}+\sigma^{2}}\right)^{2} \frac{\sigma^{2}}{n}
$$

and hence the risk is

$$
\begin{aligned}
\mathbb{E}_{\theta}\|\widehat{\theta}-\theta\|^{2} & =\sum_{j=1}^{n}\left[\left(\frac{\sigma^{2} \theta_{j}}{\sigma^{2}+c^{2}}\right)^{2}+\left(\frac{c^{2}}{c^{2}+\sigma^{2}}\right)^{2}\left(\frac{\sigma^{2}}{n}\right)\right] \\
& =\left(\frac{\sigma^{2}}{\sigma^{2}+c^{2}}\right)^{2} \sum_{j=1}^{n} \theta_{j}^{2}+\sigma^{2}\left(\frac{\sigma^{2}}{\sigma^{2}+c^{2}}\right)^{2} \\
& \leq c^{2}\left(\frac{\sigma^{2}}{\sigma^{2}+c^{2}}\right)^{2}+\sigma^{2}\left(\frac{\sigma^{2}}{\sigma^{2}+c^{2}}\right)^{2} \\
& =\frac{\sigma^{2} c^{2}}{\sigma^{2}+c^{2}} .
\end{aligned}
$$

Hence,

$$
R_{n} \leq \frac{c^{2} \sigma^{2}}{c^{2}+\sigma^{2}}
$$

for all $n$.
Lower Bound. Fix $0<\delta<1$. Let $\pi_{n}$ be a Normal prior for which $\theta_{1}, \ldots, \theta_{n}$ are IID $N\left(0, c^{2} \delta^{2} / n\right)$. Let $B\left(\pi_{n}\right)$ denote the Bayes risk. Recall that $B\left(\pi_{n}\right)$ minimizes the integrated risk $B\left(\widehat{\theta}, \pi_{n}\right)$ over all estimators. The minimum is obtained by taking $\widehat{\theta}$ to be the posterior mean which has coordinates $\widehat{\theta}_{j}=c^{2} \delta^{2} Z_{j} /\left(c^{2} \delta^{2}+\sigma^{2}\right)$ with risk

$$
R(\theta, \widehat{\theta})=\sum_{i=1}^{n}\left[\theta_{i}^{2}\left(\frac{\sigma_{n}^{2}}{\frac{c^{2} \delta^{2}}{n}+\sigma_{n}^{2}}\right)^{2}+\sigma^{2}\left(\frac{\frac{c^{2} \delta^{2}}{n}}{\frac{c^{2} \delta^{2}}{n}+\sigma_{n}^{2}}\right)^{2}\right]
$$

The Bayes risk is

$$
B\left(\pi_{n}\right)=\int R(\theta, \widehat{\theta}) d \pi_{n}(\theta)=\frac{\sigma^{2} \delta^{2} c^{2}}{\sigma^{2}+\delta^{2} c^{2}}
$$

So, for any estimator $\widehat{\theta}$,

$$
\begin{aligned}
B\left(\pi_{n}\right) & \leq B\left(\widehat{\theta}, \pi_{n}\right) \\
& =\int_{\Theta_{n}} R(\theta, \widehat{\theta}) d \pi_{n}+\int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n} \\
& \leq \sup _{\theta \in \Theta_{n}} R(\theta, \widehat{\theta})+\int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n} \\
& \leq \sup _{\theta \in \Theta_{n}} R(\theta, \widehat{\theta})+\sup _{\widehat{\theta}} \int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n}
\end{aligned}
$$

Taking the infimum over all estimators that take values in $\Theta_{n}$ yields

$$
B\left(\pi_{n}\right) \leq R_{n}+\sup _{\widehat{\theta}} \int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n}
$$

Hence,

$$
\begin{aligned}
R_{n} & \geq B\left(\pi_{n}\right)-\sup _{\hat{\theta}} \int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n} \\
& =\frac{\sigma^{2} \delta^{2} c^{2}}{\delta^{2} c^{2}+\sigma^{2}}-\sup _{\hat{\theta}} \int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n}
\end{aligned}
$$

Now, using the fact that $\|a+b\|^{2} \leq 2\left(\|a\|^{2}+\|b\|^{2}\right)$, and the CauchySchwartz inequality,

$$
\begin{aligned}
\sup _{\widehat{\theta}} \int_{\Theta_{n}^{c}} R(\theta, \widehat{\theta}) d \pi_{n} & \leq 2 \int_{\Theta_{n}^{c}}\|\theta\|^{2} d \pi_{n}+2 \sup _{\hat{\theta}} \int_{\Theta_{n}^{c}} \mathbb{E}_{\theta}\|\widehat{\theta}\|^{2} d \pi_{n} \\
& \leq 2 \sqrt{\pi_{n}\left(\Theta_{n}^{c}\right)} \sqrt{\mathbb{E}_{\pi_{n}}\left(\sum_{j} \theta_{j}^{2}\right)^{2}}+2 c^{2} \pi_{n}\left(\Theta_{n}^{c}\right)
\end{aligned}
$$

Thus,

$$
\begin{equation*}
R_{n} \geq \frac{\sigma^{2} \delta^{2} c^{2}}{\sigma^{2}+\delta^{2} c^{2}}-2 \sqrt{\pi_{n}\left(\Theta_{n}^{c}\right)} \sqrt{\mathbb{E}_{\pi_{n}}\left(\sum_{j} \theta_{j}^{2}\right)^{2}}-2 c^{2} \pi_{n}\left(\Theta_{n}^{c}\right) \tag{7.94}
\end{equation*}
$$

We now bound the last two terms in (7.94).
We shall make use of the following large deviation inequality: if $Z_{1}, \ldots, Z_{n} \sim$ $N(0,1)$ and $0<t<1$, then

$$
\mathbb{P}\left(\left|\frac{1}{n} \sum_{j}\left(Z_{j}^{2}-1\right)\right|>t\right) \leq 2 e^{-n t^{2} / 8}
$$

Let $Z_{j}=\sqrt{n} \theta_{j} /(c \delta)$ and let $t=\left(1-\delta^{2}\right) / \delta^{2}$. Then,

$$
\begin{aligned}
\pi_{n}\left(\Theta_{n}^{c}\right) & =\mathbb{P}\left(\sum_{j=1}^{n} \theta_{j}^{2}>c^{2}\right)=\mathbb{P}\left(\frac{1}{n} \sum_{j=1}^{n}\left(Z_{j}^{2}-1\right)>t\right) \\
& \leq \mathbb{P}\left(\left|\frac{1}{n} \sum_{j}\left(Z_{j}^{2}-1\right)\right|>t\right) \leq 2 e^{-n t^{2} / 8} .
\end{aligned}
$$

Next, we note that

$$
\begin{aligned}
\mathbb{E}_{\pi_{n}}\left(\sum_{j} \theta_{j}^{2}\right)^{2} & =\sum_{i=1}^{n} \mathbb{E}_{\pi_{n}}\left(\theta_{i}^{4}\right)+\sum_{i=1}^{n} \sum_{j \neq i}^{n} \mathbb{E}_{\pi_{n}}\left(\theta_{i}^{2}\right) \mathbb{E}_{\pi_{n}}\left(\theta_{j}^{2}\right) \\
& =\frac{c^{4} \delta^{4} \mathbb{E}\left(Z_{1}^{4}\right)}{n}+\binom{n}{2} \frac{c^{4} \delta^{4}}{n^{2}}=O(1)
\end{aligned}
$$

Therefore, from (7.94),

$$
R_{n} \geq \frac{\sigma^{2} \delta^{2} c^{2}}{\sigma^{2}+\delta^{2} c^{2}}-2 \sqrt{2} e^{-n t^{2} / 16} O(1)-2 c^{2} e^{-n t^{2} / 8}
$$

Hence,

$$
\liminf _{n \rightarrow \infty} R_{n} \geq \frac{\sigma^{2} \delta^{2} c^{2}}{\sigma^{2}+\delta^{2} c^{2}}
$$

The conclusion follows by letting $\delta \uparrow 1$.
Proof of Theorem 7.74. Let

$$
a=\frac{\sigma_{n}}{n^{1 / 4}}\left(\log \left(1+\epsilon^{2}\right)\right)^{1 / 4}
$$

and define

$$
\Omega=\left\{\theta=\left(\theta_{1}, \ldots, \theta_{n}\right):\left|\theta_{i}\right|=a, \quad i=1, \ldots, n\right\}
$$

Note that $\Omega$ contains $2^{n}$ elements. Let $f_{\theta}$ denote the density of a multivariate Normal with mean $\theta$ and covariance $\sigma_{n}^{2} I$ where $I$ is the identity matrix. Define the mixture

$$
q(y)=\frac{1}{2^{n}} \sum_{\theta \in \Omega} f_{\theta}(y)
$$

Let $f_{0}$ denote the density of a multivariate Normal with mean $(0, \ldots, 0)$ and covariance $\sigma_{n}^{2} I$. Then,

$$
\begin{aligned}
\int\left|f_{0}(x)-g(x)\right| d x & =\int \frac{\left|f_{0}(x)-g(x)\right|}{\sqrt{f_{0}(x)}} \sqrt{f_{0}(x)} d x \\
& \leq \sqrt{\int \frac{\left(f_{0}(x)-g(x)\right)^{2}}{f_{0}(x)}} d x \\
& =\sqrt{\int \frac{g^{2}(x)}{f_{0}(x)} d x-1} .
\end{aligned}
$$

Now,

$$
\begin{aligned}
& \int \frac{q^{2}(x)}{f_{0}(x)} d x=\int\left(\frac{q(x)}{f_{0}(x)}\right)^{2} f_{0}(x) d x=\mathbb{E}_{0}\left(\frac{q(x)}{f_{0}(x)}\right)^{2} \\
& \quad=\left(\frac{1}{2^{n}}\right)^{2} \sum_{\theta, \nu \in \Omega} \mathbb{E}_{0}\left(\frac{f_{\theta}(x) f_{\nu}(x)}{f_{0}^{2}(x)}\right) \\
& \quad=\left(\frac{1}{2^{n}}\right)^{2} \sum_{\theta, \nu \in \Omega} \exp \left\{-\frac{1}{2 \sigma_{n}^{2}}\left(\|\theta\|^{2}+\|\nu\|^{2}\right)\right\} \mathbb{E}_{0}\left(\exp \left\{\epsilon^{T}(\theta+\nu) / \sigma_{n}^{2}\right\}\right) \\
& \quad=\left(\frac{1}{2^{n}}\right)^{2} \sum_{\theta, \nu \in \Omega} \exp \left\{-\frac{1}{2 \sigma_{n}^{2}}\left(\|\theta\|^{2}+\|\nu\|^{2}\right)\right\} \exp \left\{\sum_{i=1}^{n}\left(\theta_{i}+\nu_{i}\right)^{2} /\left(2 \sigma_{n}^{2}\right)\right\} \\
& \quad=\left(\frac{1}{2^{n}}\right)^{2} \sum_{\theta, \nu \in \Omega} \exp \left\{\frac{\langle\theta, \nu\rangle}{\sigma_{n}^{2}}\right\}
\end{aligned}
$$

The latter is equal to the mean of $\exp \left(\langle\theta, \nu\rangle / \sigma_{n}^{2}\right)$ when drawing two vectors $\theta$ and $\nu$ at random from $\Omega$. And this, in turn, is equal to

$$
\mathbb{E} \exp \left\{\frac{a^{2} \sum_{i=1}^{n} E_{i}}{\sigma_{n}^{2}}\right\}
$$

where $E_{1}, \ldots, E_{n}$ are independent and $\mathbb{P}\left(E_{i}=1\right)=\mathbb{P}\left(E_{i}=-1\right)=1 / 2$. Moreover,

$$
\begin{aligned}
\mathbb{E} \exp \left\{\frac{a^{2} \sum_{i=1}^{n} E_{i}}{\sigma_{n}^{2}}\right\} & =\prod_{i=1}^{n} \mathbb{E} \exp \left\{\frac{a^{2} E_{i}}{\sigma_{n}^{2}}\right\} \\
& =\left(\mathbb{E} \exp \left\{\frac{a^{2} E_{1}}{\sigma_{n}^{2}}\right\}\right)^{n} \\
& =\left(\cosh \left(\frac{a^{2}}{\sigma_{n}^{2}}\right)\right)^{n}
\end{aligned}
$$

where $\cosh (y)=\left(e^{y}+e^{-y}\right) / 2$. Thus,

$$
\int \frac{q^{2}(x)}{f_{0}(x)} d x=\left(\cosh \left(\frac{a^{2}}{\sigma_{n}^{2}}\right)\right)^{n} \leq e^{a^{4} n / \sigma_{n}^{4}}
$$

where we have used the fact that $\cosh (y) \leq e^{y^{2}}$. Thus,

$$
\int\left|f_{0}(x)-q(x)\right| d x \leq \sqrt{e^{a^{4} n / \sigma_{n}^{4}}-1}=\epsilon
$$

So, if $Q$ denotes the probability measure with density $q$, we have, for any event $A$,

$$
\begin{align*}
Q(A) & =\int_{A} q(x) d x=\int_{A} f_{0}(x) d x+\int_{A}\left(q(x)-f_{0}(x)\right) d x \\
& \geq \mathbb{P}_{0}(A)-\int_{A}\left|q(x)-f_{0}(x)\right| d x \geq \mathbb{P}_{0}(A)-\epsilon \tag{7.95}
\end{align*}
$$

Define two events, $A=\left\{(0, \ldots, 0) \in \mathcal{B}_{n}\right\}$ and $B=\left\{\Omega \bigcap \mathcal{B}_{n} \neq \emptyset\right\}$. Every $\theta \in \Omega$ has norm

$$
\|\theta\|=\sqrt{n a^{2}}=\sigma_{n} n^{1 / 4}\left(\log \left(1+\epsilon^{2}\right)\right)^{1 / 4} \equiv c_{n} .
$$

Hence, $A \bigcap B \subset\left\{s_{n} \geq c_{n}\right\}$. Since $\mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right) \geq 1-\alpha$ for all $\theta$, it follows that $\mathbb{P}_{\theta}(B) \geq 1-\alpha$ for all $\theta \in \Omega$. Hence, $Q(B) \geq 1-\alpha$. From (7.95),

$$
\begin{aligned}
\mathbb{P}_{0}\left(s_{n} \geq c_{n}\right) & \geq \mathbb{P}_{0}(A \bigcap B) \geq Q(A \bigcap B)-\epsilon \\
& =Q(A)+Q(B)-Q(A \bigcup B)-\epsilon \\
& \geq Q(A)+Q(B)-1-\epsilon \\
& \geq Q(A)+(1-\alpha)-1-\epsilon \\
& \geq \mathbb{P}_{0}(A)+(1-\alpha)-1-2 \epsilon \\
& \geq(1-\alpha)+(1-\alpha)-1-2 \epsilon \\
& =1-2 \alpha-2 \epsilon .
\end{aligned}
$$

So, $\mathbb{E}_{0}\left(s_{n}\right) \geq(1-2 \alpha-2 \epsilon) c_{n}$. It is easy to see that the same argument can be used for any $\theta \in \mathbb{R}^{n}$ and hence $\mathbb{E}_{\theta}\left(s_{n}\right) \geq(1-2 \alpha-2 \epsilon) c_{n}$ for every $\theta \in \mathbb{R}^{n}$.

Proof of Theorem 7.77. Let $a=\sigma_{n} z_{\alpha+2 \epsilon}$ where $0<\epsilon<(1 / 2)(1 / 2-\alpha)$ and define

$$
\Omega=\left\{\theta=\left(\theta_{1}, \ldots, \theta_{n}\right):\left|\theta_{i}\right|=a, \quad i=1, \ldots, n\right\}
$$

Define the loss function $L=L(\widehat{\theta}, \theta)=\sum_{i=1}^{n} I\left(\left|\widehat{\theta}_{i}-\theta_{i}\right| \geq a\right)$. Let $\pi$ be the uniform prior on $\Omega$. The posterior mass function over $\Omega$ is $p(\theta \mid y)=\prod_{i=1}^{n} p\left(\theta_{i} \mid y_{i}\right)$ where

$$
p\left(\theta_{i} \mid y_{i}\right)=\frac{e^{2 a y_{i} / \sigma_{n}^{2}}}{1+e^{2 a y_{i} / \sigma_{n}^{2}}} I\left(\theta_{i}=a\right)+\frac{1}{1+e^{2 a y_{i} / \sigma_{n}^{2}}} I\left(\theta_{i}=-a\right) .
$$

The posterior risk is

$$
\mathbb{E}(L(\widehat{\theta}, \theta) \mid y)=\sum_{i=1}^{n} \mathbb{P}\left(\left|\widehat{\theta}_{i}-\theta_{i}\right| \geq a \mid y_{i}\right)
$$

which is minimized by taking $\widehat{\theta}_{i}=a$ if $y_{i} \geq 0$ and $\widehat{\theta}_{i}=-a$ if $y_{i}<0$. The risk of this estimator is

$$
\begin{aligned}
& \sum_{i=1}^{n}\left(\mathbb{P}\left(Y_{i}<0 \mid \theta_{i}=a\right) I\left(\theta_{i}=a\right)+\mathbb{P}\left(Y_{i}>0 \mid \theta_{i}=-a\right) I\left(\theta_{i}=-a\right)\right) \\
& \quad=n \Phi\left(-a / \sigma_{n}\right)=n(\alpha+2 \epsilon)
\end{aligned}
$$

Since this risk is constant, it is the minimax risk. Therefore,

$$
\begin{aligned}
\inf _{\widehat{\theta}} \sup _{\theta \in \mathbb{R}^{n}} \sum_{i=1}^{n} \mathbb{P}_{\theta}\left(\left|\widehat{\theta}_{i}-\theta_{i}\right| \geq a\right) & \geq \inf _{\widehat{\theta}} \sup _{\theta \in \Omega} \sum_{i=1}^{n} \mathbb{P}_{\theta}\left(\left|\widehat{\theta}_{i}-\theta_{i}\right| \geq a\right) \\
& =n(\alpha+2 \epsilon)
\end{aligned}
$$

Let $\gamma=\epsilon /(1-\alpha-\epsilon)$. Given any estimator $\widehat{\theta}$,

$$
\gamma n \mathbb{P}_{\theta}(L<\gamma n)+n \mathbb{P}_{\theta}(L \geq \gamma n) \geq L
$$

and so

$$
\sup _{\theta}\left(\gamma n \mathbb{P}_{\theta}(L<\gamma n)+n \mathbb{P}_{\theta}(L \geq \gamma n)\right) \geq \sup _{\theta} \mathbb{E}_{\theta}(L) \geq n(\alpha+2 \epsilon)
$$

This inequality, together with the fact that $\mathbb{P}_{\theta}(L<\gamma n)+\mathbb{P}_{\theta}(L \geq \gamma n)=1$ implies that

$$
\sup _{\theta} \mathbb{P}_{\theta}(L \geq \gamma n) \geq \alpha+\epsilon
$$

Thus,

$$
\sup _{\theta} \mathbb{P}_{\theta}\left(\|\widehat{\theta}-\theta\|^{2} \geq \gamma n a^{2}\right) \geq \sup _{\theta} \mathbb{P}_{\theta}(L \geq \gamma n) \geq \alpha+\epsilon .
$$

Therefore,

$$
\begin{aligned}
\sup _{\theta} \mathbb{P}_{\theta}\left(s_{n}^{2} \geq \gamma n a^{2}\right) & \geq \sup _{\theta} \mathbb{P}_{\theta}\left(s_{n}^{2} \geq\|\widehat{\theta}-\theta\|^{2} \geq \gamma n a^{2}\right) \\
& =\sup _{\theta} \mathbb{P}_{\theta}\left(s_{n}^{2} \geq\|\widehat{\theta}-\theta\|^{2}\right)+\sup _{\theta} \mathbb{P}_{\theta}\left(\|\widehat{\theta}-\theta\|^{2} \geq \gamma n a^{2}\right)-1 \\
& \geq \alpha+\epsilon+1-\alpha-1=\epsilon .
\end{aligned}
$$

Thus, $\sup _{\theta} E_{\theta}\left(s_{n}\right) \geq \epsilon a \sqrt{\gamma n}$.

### 7.14 Exercises

1. Let $\theta_{i}=1 / i^{2}$ for $i=1, \ldots, n$. Take $n=1000$. Let $Z_{i} \sim N\left(\theta_{i}, 1\right)$ for $i=1, \ldots, n$. Compute the risk of the mLE. Compute the risk of the estimator $\widetilde{\theta}=\left(b Z_{1}, b Z_{2}, \ldots, b Z_{n}\right)$. Plot this risk as a function of $b$. Find the optimal value $b_{*}$. Now conduct a simulation. For each run of the simulation, find the (modified) James-Stein estimator $\widehat{b} Z$ where

$$
\widehat{b}=\left[1-\frac{n}{\sum_{i} Z_{i}^{2}}\right]^{+}
$$

You will get one $\widehat{b}$ for each simulation. Compare the simulated values of $\widehat{b}$ to $b_{*}$. Also, compare the risk of the MLE and the James-Stein estimator (the latter obtained by simulation) to the Pinsker bound.
2. For the Normal means problem, consider the following curved soft threshold estimator:

$$
\widehat{\theta}_{i}= \begin{cases}-\left(Z_{i}+\lambda\right)^{2} & Z_{i}<-\lambda \\ 0 & -\lambda \leq Z_{i} \leq \lambda \\ \left(Z_{i}-\lambda\right)^{2} & Z_{i}>\lambda\end{cases}
$$

where $\lambda>0$ is some fixed constant.
(a) Find the risk of this estimator. Hint: $R=\mathbb{E}$ (SURE).
(b) Consider problem (1). Use your estimator from (2a) with $\lambda$ chosen from the data using SURE. Compare the risk to the risk of the JamesStein estimator. Now repeat the comparison for

$$
\theta=(\overbrace{10, \ldots, 10}^{10 \text { times }}, \overbrace{0, \ldots, 0}^{990 \text { times }}) .
$$

3. Let $J=J_{n}$ be such that $J_{n} \rightarrow \infty$ and $n \rightarrow \infty$. Let

$$
\widehat{\sigma}^{2}=\frac{n}{J} \sum_{i=n-J+1}^{n} Z_{i}^{2}
$$

where $Z_{i} \sim N\left(\theta_{i}, \sigma^{2} / n\right)$. Show that if $\theta=\left(\theta_{1}, \theta_{2}, \ldots\right)$ belongs to a Sobolev body of order $m>1 / 2$ then $\widehat{\sigma}^{2}$ is a uniformly consistent estimator of $\sigma^{2}$ in the Normal means model.
4. Prove Stein's lemma: if $X \sim N\left(\mu, \sigma^{2}\right)$ then $\mathbb{E}(g(X)(X-\mu))=\sigma^{2} \mathbb{E} g^{\prime}(X)$.
5. Verify equation 7.22 .
6. Show that the hard threshold estimator defined in (7.23) is not weakly differentiable.
7. Compute the risk functions for the soft threshold estimator (7.21) and the hard threshold estimator (7.23).
8. Generate $Z_{i} \sim N\left(\theta_{i}, 1\right), i=1, \ldots, 100$, where $\theta_{i}=1 / i$. Compute a 95 percent confidence ball using: (i) the $\chi^{2}$ confidence ball, (ii) the Baraud method, (iii) the pivotal method. Repeat 1000 times and compare the radii of the balls.
9. Let $\|a-b\|_{\infty}=\sup _{j}\left|a_{j}-b_{j}\right|$. Construct a confidence set $B_{n}$ of the form $B_{n}=\left\{\theta \in \mathbb{R}^{n}:\left\|\theta-Z^{n}\right\|_{\infty} \leq c_{n}\right\}$ such that $\mathbb{P}_{\theta}\left(\theta \in B_{n}\right) \geq 1-\alpha$ for all $\theta \in \mathbb{R}^{n}$ under model (7.1) with $\sigma_{n}=\sigma / \sqrt{n}$. Find the expected diameter of your confidence set.
10. Consider Example 7.24. Define

$$
\delta=\max _{S \in \mathcal{S}} \sup _{\theta \in \mathbb{R}^{n}}\left|\widehat{R}_{S}-R\left(\widehat{\theta}_{S}, \theta\right)\right|
$$

Try to bound $\delta$ in the following three cases: (i) $\mathcal{S}$ consists of a single model $S$; (ii) nested model selection; (iii) all subsets selection.
11. Consider Example 7.24. Another method for choosing a model is to use penalized likelihood. In particular, some well-known penalization model selection methods are AIC (Akaike's Information Criterion), Akaike (1973), Mallows' $C_{p}$, Mallows (1973), and BIC (Bayesian Information Criterion), Schwarz (1978). In the Normal means model, minimizing
sure, AIC and $C_{p}$ are equivalent. But BIC leads to a different model selection procedure. Specifically,

$$
\mathrm{BIC}_{B}=\ell_{B}-\frac{|B|}{2} \log n
$$

where $\ell_{B}$ is the log-likelihood of the submodel $B$ evaluated at its maximum likelihood estimator. Find an explicit expression for $\mathrm{BIC}_{B}$. Suppose we choose $B$ by maximizing $\operatorname{BIC}_{B}$ over $\mathcal{B}$. Investigate the properties of this model selection procedure and compare it to selecting a model by minimizing SURE. In particular, compare the risk of the resulting estimators. Also, assuming there is a "true" submodel (that is, $\theta_{i} \neq 0$ if and only if $i \in B$ ), compare the probability of selecting the true submodel under each procedure. In general, estimating $\theta$ accurately and finding the true submodel are not the same. See Wasserman (2000).
12. By approximating the noncentral $\chi^{2}$ with a Normal, find a large sample approximation to for $\rho_{0}$ and $\rho_{n}$ in Example 7.69. Then prove equation (7.70).

## 8

## Nonparametric Inference Using Orthogonal Functions

### 8.1 Introduction

In this chapter we use orthogonal function methods for nonparametric inference. Specifically, we use an orthogonal basis to convert regression and density estimation into a Normal means problem and then we construct estimates and confidence sets using the theory from Chapter 7. In the regression case, the resulting estimators are linear smoothers and thus are a special case of the estimators described in Section 5.2. We discuss another approach to orthogonal function regression based on wavelets in the next chapter.

### 8.2 Nonparametric Regression

The particular version of orthogonal function regression that we consider here was developed by Beran (2000) and Beran and Dümbgen (1998). They call the method react, which stands for Risk Estimation and Adaptation after Coordinate Transformation. Similar ideas have been developed by Efromovich (1999). In fact, the basic idea is quite old; see Cenčov (1962), for example.

Suppose we observe

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\sigma \epsilon_{i} \tag{8.1}
\end{equation*}
$$

where $\epsilon_{i} \sim N(0,1)$ are IID. For now, we assume a regular design meaning that $x_{i}=i / n, i=1, \ldots, n$.

Let $\phi_{1}, \phi_{2}, \ldots$ be an orthonormal basis for $[0,1]$. In our examples we will often use the cosine basis:

$$
\begin{equation*}
\phi_{1}(x) \equiv 1, \quad \phi_{j}(x)=\sqrt{2} \cos ((j-1) \pi x), j \geq 2 \tag{8.2}
\end{equation*}
$$

Expand $r$ as

$$
\begin{equation*}
r(x)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(x) \tag{8.3}
\end{equation*}
$$

where $\theta_{j}=\int_{0}^{1} \phi_{j}(x) r(x) d x$.
First, we approximate $r$ by

$$
r_{n}(x)=\sum_{j=1}^{n} \theta_{j} \phi_{j}(x)
$$

which is the projection of $r$ onto the span of $\left\{\phi_{1}, \ldots, \phi_{n}\right\} .{ }^{1}$ This introduces an integrated squared bias of size

$$
B_{n}(\theta)=\int_{0}^{1}\left(r(x)-r_{n}(x)\right)^{2} d x=\sum_{j=n+1}^{\infty} \theta_{j}^{2}
$$

If $r$ is smooth, this bias is quite small.
8.4 Lemma. Let $\Theta(m, c)$ be a Sobolev ellipsoid. ${ }^{2}$ Then,

$$
\begin{equation*}
\sup _{\theta \in \Theta(m, c)} B_{n}(\theta)=O\left(\frac{1}{n^{2 m}}\right) \tag{8.5}
\end{equation*}
$$

In particular, if $m>1 / 2$ then $\sup _{\theta \in \Theta(m, c)} B_{n}(\theta)=o(1 / n)$.
Hence this bias is negligible and we shall ignore it for the rest of the chapter. More precisely, we will focus on estimating $r_{n}$ rather than $r$. Our next task is to estimate the $\theta=\left(\theta_{1}, \ldots, \theta_{n}\right)$. Let

$$
\begin{equation*}
Z_{j}=\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(x_{i}\right), \quad j=1, \ldots \tag{8.6}
\end{equation*}
$$

[^18]As we saw in equation (7.15) of Chapter 7,

$$
\begin{equation*}
Z_{j} \approx N\left(\theta_{j}, \frac{\sigma^{2}}{n}\right) \tag{8.7}
\end{equation*}
$$

We know from the previous chapter that the mLE $Z=\left(Z_{1}, \ldots, Z_{n}\right)$ has large risk. One possibility for improving on the MLE is to use the James-Stein estimator $\widehat{\theta}_{J S}$ defined in (7.41). We can think of the James-Stein estimator as the estimator that minimizes the estimated risk over all estimators of the form $\left(b Z_{1}, \ldots, b Z_{n}\right)$. REACT generalizes this idea by minimizing the risk over a larger class of estimators, called modulators.

A modulator is a vector $b=\left(b_{1}, \ldots, b_{n}\right)$ such that $0 \leq b_{j} \leq 1, j=1, \ldots, n$. A modulation estimator is an estimator of the form

$$
\begin{equation*}
\widehat{\theta}=b Z=\left(b_{1} Z_{1}, b_{2} Z_{2}, \ldots, b_{n} Z_{n}\right) \tag{8.8}
\end{equation*}
$$

A constant modulator is a modulator of the form $(b, \ldots, b)$. A nested subset selection modulator is a modulator of the form

$$
b=(1, \ldots, 1,0, \ldots, 0)
$$

A monotone modulator is a modulator of the form

$$
1 \geq b_{1} \geq \cdots \geq b_{n} \geq 0
$$

The set of constant modulators is denoted by $\mathcal{M}_{\text {CONS }}$, the set of nested subset modulators is denoted by $\mathcal{M}_{\mathrm{NSS}}$ and the set of monotone modulators is denoted by $\mathcal{M}_{\text {MON }}$.

Given a modulator $b=\left(b_{1}, \ldots, b_{n}\right)$, the function estimator is

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{j=1}^{n} \widehat{\theta}_{j} \phi_{j}(x)=\sum_{j=1}^{n} b_{j} Z_{j} \phi_{j}(x) \tag{8.9}
\end{equation*}
$$

Observe that

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x) \tag{8.10}
\end{equation*}
$$

where

$$
\begin{equation*}
\ell_{i}(x)=\frac{1}{n} \sum_{j=1}^{n} b_{j} \phi_{j}(x) \phi_{j}\left(x_{i}\right) \tag{8.11}
\end{equation*}
$$

Hence, $\widehat{r}_{n}$ is a linear smoother as described in Section 5.2.
Modulators shrink the $Z_{j}$ s towards 0 and, as we saw in the last chapter, shrinking tends to smooth the function. Thus, choosing the amount of shrinkage corresponds to the problem of choosing a bandwidth that we faced in

Chapter 5. We shall address the problem using Stein's unbiased risk estimator (Section 7.4) instead of cross-validation.

Let

$$
R(b)=\mathbb{E}_{\theta}\left[\sum_{j=1}^{n}\left(b_{j} Z_{j}-\theta_{j}\right)^{2}\right]
$$

denote the risk of the estimator $\widehat{\theta}=\left(b_{1} Z_{1}, \ldots, b_{n} Z_{n}\right)$. The idea of REACT is to estimate the risk $R(b)$ and choose $\widehat{b}$ to minimize the estimated risk over a class of modulators $\mathcal{M}$. Minimizing over $\mathcal{M}_{\text {CONS }}$ yields the James-Stein estimator, so REACT is a generalization of James-Stein estimation.

To proceed, we need to estimate $\sigma$. Any of the methods discussed in Chapter 5 can be used. Another estimator, well-suited for the present framework, is

$$
\begin{equation*}
\widehat{\sigma}^{2}=\frac{1}{n-J_{n}} \sum_{i=n-J_{n}+1}^{n} Z_{i}^{2} . \tag{8.12}
\end{equation*}
$$

This estimator is consistent as long as $J_{n} \rightarrow \infty$ and $n-J_{n} \rightarrow \infty$ as $n \rightarrow \infty$. As a default value, $J_{n}=n / 4$ is not unreasonable. The intuition is that if $r$ is smooth then we expect $\theta_{j} \approx 0$ for large $j$, and hence $Z_{j}^{2}=\left(\theta_{j}^{2}+\sigma \epsilon_{j} / \sqrt{n}\right)^{2} \approx$ $\left(\sigma \epsilon_{j} / \sqrt{n}\right)^{2}=\sigma^{2} \epsilon_{j}^{2} / n$. Therefore,

$$
\mathbb{E}\left(\widehat{\sigma}^{2}\right)=\frac{1}{n-J_{n}} \sum_{n-J_{n}+1}^{n} \mathbb{E}\left(Z_{i}^{2}\right) \approx \frac{1}{n-J_{n}} \sum_{n-J_{n}+1}^{n} \frac{\sigma^{2}}{n} \mathbb{E}\left(\epsilon_{i}^{2}\right)=\sigma^{2} .
$$

Now we can estimate the risk function.
8.13 Theorem. The risk of a modulator $b$ is

$$
\begin{equation*}
R(b)=\sum_{j=1}^{n} \theta_{j}^{2}\left(1-b_{j}\right)^{2}+\frac{\sigma^{2}}{n} \sum_{j=1}^{n} b_{j}^{2} . \tag{8.14}
\end{equation*}
$$

The (modified) ${ }^{3}$ SURE estimator of $R(b)$ is

$$
\begin{equation*}
\widehat{R}(b)=\sum_{j=1}^{n}\left(Z_{j}^{2}-\frac{\widehat{\sigma}^{2}}{n}\right)_{+}\left(1-b_{j}\right)^{2}+\frac{\widehat{\sigma}^{2}}{n} \sum_{j=1}^{n} b_{j}^{2} \tag{8.15}
\end{equation*}
$$

where $\widehat{\sigma}^{2}$ is a consistent estimate of $\sigma^{2}$ such as (8.12).

[^19]
### 8.16 Definition. Let $\mathcal{M}$ be a set of modulators. The modulation

 estimator of $\theta$ is $\widehat{\theta}=\left(\widehat{b}_{1} Z_{1}, \ldots, \widehat{b}_{n} Z_{n}\right)$ where $\widehat{b}=\left(\widehat{b}_{1} \ldots, \widehat{b}_{n}\right)$ minimizes $\widehat{R}(b)$ over $\mathcal{M}$. The REACT function estimator is$$
\widehat{r}_{n}(x)=\sum_{j=1}^{n} \widehat{\theta}_{j} \phi_{j}(x)=\sum_{j=1}^{n} \widehat{b}_{j} Z_{j} \phi_{j}(x) .
$$

For a fixed $b$, we expect that $\widehat{R}(b)$ approximates $R(b)$. But for the REACT estimator we require more: we want $\widehat{R}(b)$ to approximate $R(b)$ uniformly for $b \in \mathcal{M}$. If so, then $\inf _{b \in \mathcal{M}} \widehat{R}(b) \approx \inf _{b \in \mathcal{M}} R(b)$ and the $b$ that minimizes $\widehat{R}(b)$ should be nearly as good as the $b$ that minimizes $R(b)$. This motivates the next result.
8.17 Theorem (Beran and Dümbgen, 1998). Let $\mathcal{M}$ be one of $\mathcal{M}_{\text {CONS }}, \mathcal{M}_{\text {NSS }}$ or $\mathcal{M}_{\mathrm{MON}}$. Let $R(b)$ denote the true risk of the estimator $\left(b_{1} Z_{1}, \ldots, b_{n} Z_{n}\right)$. Let $b^{*}$ minimize $R(b)$ over $\mathcal{M}$ and let $\widehat{b}$ minimize $\widehat{R}(b)$ over $\mathcal{M}$. Then

$$
\left|R(\widehat{b})-R\left(b^{*}\right)\right| \rightarrow 0
$$

as $n \rightarrow \infty$. For $\mathcal{M}=\mathcal{M}_{\text {CONS }}$ or $\mathcal{M}=\mathcal{M}_{\mathrm{MON}}$, the estimator $\widehat{\theta}=\left(\widehat{b}_{1} Z_{1}, \ldots\right.$, $\widehat{b}_{n} Z_{n}$ ) achieves the Pinsker bound (7.29).

To implement this method, we need to find $\widehat{b}$ to minimize $\widehat{R}(b)$. The minimum of $\widehat{R}(b)$ over $\mathcal{M}_{\text {CONS }}$ is the James-Stein estimator. To minimize $\widehat{R}(b)$ over $\mathcal{M}_{\text {NSS }}$, we compute $\widehat{R}(b)$ for every modulator of the form $(1,1 \ldots, 1,0, \ldots, 0)$ and then the minimum is found. In other words, we find $\widehat{J}$ to minimize

$$
\begin{equation*}
\widehat{R}(J)=\frac{J \widehat{\sigma}^{2}}{n}+\sum_{j=J+1}^{n}\left(Z_{j}^{2}-\frac{\widehat{\sigma}^{2}}{n}\right)_{+} \tag{8.18}
\end{equation*}
$$

and set $\widehat{r}_{n}(x)=\sum_{j=1}^{\widehat{J}} Z_{j} \phi_{j}(x)$. It is a good idea to plot the estimated risk as a function of $J$. To minimize $\widehat{R}(b)$ over $\mathcal{M}_{\text {MON }}$, note that $\widehat{R}(b)$ can be written as

$$
\begin{equation*}
\widehat{R}(b)=\sum_{i=1}^{n}\left(b_{i}-g_{i}\right)^{2} Z_{i}^{2}+\frac{\widehat{\sigma}^{2}}{n} \sum_{i=1}^{n} g_{i} \tag{8.19}
\end{equation*}
$$

where $g_{i}=\left(Z_{i}^{2}-\left(\hat{\sigma}^{2} / n\right)\right) / Z_{i}^{2}$. So it suffices to minimize

$$
\sum_{i=1}^{n}\left(b_{i}-g_{i}\right)^{2} Z_{i}^{2}
$$

subject to $b_{1} \geq \cdots \geq b_{n}$. This is simply a weighted least squares problem subject to a monotonicity constraint. There is a well-known algorithm called the pooled-adjacent-violators (PAV) algorithm for performing this minimization; See Robertson et al. (1988).

Usually, monotone modulators lead to estimates that are close to the NSS modulators and the latter are very easy to implement. Thus, as a default, the NSS method is reasonable. At this point, we can summarize the whole REACT procedure.

## Summary of REACT

1. Let $Z_{j}=n^{-1} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(x_{i}\right)$ for $j=1, \ldots, n$.
2. Find $\widehat{J}$ to minimize the risk estimator $\widehat{R}(J)$ given by equation (8.18).
3. Let

$$
\widehat{r}_{n}(x)=\sum_{j=1}^{\widehat{J}} Z_{j} \phi_{j}(x) .
$$

8.20 Example (Doppler function). Recall that the Doppler function from

Example 5.63 is

$$
r(x)=\sqrt{x(1-x)} \sin \left(\frac{2.1 \pi}{x+.05}\right) .
$$

The top left panel in Figure 8.1 shows the true function. The top right panel shows 1000 data points. The data were simulated from the model $Y_{i}=r(i / n)+$ $\sigma \epsilon_{i}$ with $\sigma=0.1$ and $\epsilon_{i} \sim N(0,1)$. The bottom left panel shows the estimated risk for the NSS modulator as a function of the number of terms in the fit. The risk was minimized by using the modulator:

$$
b=(\underbrace{1, \ldots, 1}_{187}, \underbrace{0, \ldots, 0}_{813}) .
$$

The bottom right panel shows the REACT fit. Compare with Figure 5.6.
8.21 Example (CMB data). Let us compare REACT to local smoothing for the CMB data from Example 4.4. The estimated risk (for NSS) is minimized by using $J=6$ basis functions. The fit is shown in Figure 8.2. and is similar to the fits obtained in Chapter 5. (We are ignoring the fact that the variance is not constant.) The plot of the risk reveals that there is another local minimum around $J=40$. The bottom right plot shows the fit using 40 basis functions. This fit appears to undersmooth.





FIGURE 8.1. Doppler test function. Top left: true function. Top right: 1000 data points. Bottom left: estimated risk as a function of the number of terms in the fit. Bottom right: final react fit.

There are several ways to construct a confidence set for $r$. We begin with confidence balls. First, construct a confidence ball $\mathcal{B}_{n}$ for $\theta=\left(\theta_{1}, \ldots, \theta_{n}\right)$ using any of the methods in Section 7.8. Then define

$$
\begin{equation*}
\mathcal{C}_{n}=\left\{r=\sum_{j=1}^{n} \theta_{j} \phi_{j}:\left(\theta_{1}, \ldots, \theta_{n}\right) \in \mathcal{B}_{n}\right\} . \tag{8.22}
\end{equation*}
$$

It follows that $\mathcal{C}_{n}$ is a confidence ball for $r_{n}$. If we use the pivotal method from Section 7.8 we get the following.
8.23 Theorem (Beran and Dümbgen, 1998). Let $\widehat{\theta}$ be the MON or NSS estimator and let $\widehat{\sigma}^{2}$ be the estimator of $\sigma$ defined in (8.12). Let

$$
\begin{equation*}
\mathcal{B}_{n}=\left\{\theta=\left(\theta_{1}, \ldots, \theta_{n}\right): \sum_{j=1}^{n}\left(\theta_{j}-\widehat{\theta}_{j}\right)^{2} \leq s_{n}^{2}\right\} \tag{8.24}
\end{equation*}
$$

where

$$
\begin{aligned}
s_{n}^{2}= & \widehat{R}(\widehat{b})+\frac{\widehat{\tau} z_{\alpha}}{\sqrt{n}} \\
\widehat{\tau}^{2}= & \frac{2 \widehat{\sigma}^{4}}{n} \sum_{j}\left(\left(2 \widehat{b}_{j}-1\right)\left(1-c_{j}\right)\right)^{2} \\
& +4 \widehat{\sigma}^{2} \sum_{j}\left(Z_{j}^{2}-\frac{\widehat{\sigma}^{2}}{n}\right)\left(\left(1-\widehat{b}_{j}\right)+\left(2 \widehat{b}_{j}-1\right) c_{j}\right)^{2}
\end{aligned}
$$



FIGURE 8.2. CMB data using react. Top left: NSS fit using $J=6$ basis functions. Top right: estimated risk. Bottom left: NSS fit using $J=40$ basis functions.
and

$$
c_{j}= \begin{cases}0 & \text { if } j \leq n-J \\ 1 / J & \text { if } j>n-J .\end{cases}
$$

Then, for any $c>0$ and $m>1 / 2$,

$$
\lim _{n \rightarrow \infty} \sup _{\theta \in \Theta(m, c)}\left|\mathbb{P}_{\theta}\left(\theta \in \mathcal{B}_{n}\right)-(1-\alpha)\right|=0 .
$$

To construct confidence bands, we use the fact that $\widehat{r}_{n}$ is a linear smoother and we can then use the method from Section 5.7. The band is given by (5.99), namely,

$$
\begin{equation*}
I(x)=\left(\widehat{r}_{n}(x)-c \widehat{\sigma}\|\ell(x)\|, \widehat{r}_{n}(x)+c \widehat{\sigma}\|\ell(x)\|\right) \tag{8.25}
\end{equation*}
$$

where

$$
\begin{equation*}
\|\ell(x)\|^{2} \approx \frac{1}{n} \sum_{j=1}^{n} b_{j}^{2} \phi_{j}^{2}(x) \tag{8.26}
\end{equation*}
$$

and $c$ is from equation (5.102).

### 8.3 Irregular Designs

So far we have assumed a regular design $x_{i}=i / n$. Now let us relax this assumption and deal with an irregular design. There are several ways to
handle this case. The simplest is to use a basis $\left\{\phi_{1}, \ldots, \phi_{n}\right\}$ that is orthogonal with respect to the design points $x_{1}, \ldots, x_{n}$. That is, we choose a basis for $L_{2}\left(P_{n}\right)$ where $P_{n}=n^{-1} \sum_{i=1}^{n} \delta_{i}$ and $\delta_{i}$ is a point mass at $x_{i}$. This requires that

$$
\left\|\phi_{j}^{2}\right\|=1, \quad j=1, \ldots, n
$$

and

$$
\left\langle\phi_{j}, \phi_{k}\right\rangle=0, \quad 1 \leq j<k \leq n
$$

where

$$
\langle f, g\rangle=\int f(x) g(x) d P_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} f\left(x_{i}\right) g\left(x_{i}\right)
$$

and

$$
\|f\|^{2}=\int f^{2}(x) d P_{n}(x)=\frac{1}{n} \sum_{i=1}^{n} f^{2}\left(x_{i}\right)
$$

We can construct such a basis by Gram-Schmidt orthogonalization as follows.
Let $g_{1}, \ldots, g_{n}$ be any convenient orthonormal basis for $\mathbb{R}^{n}$. Let

$$
\phi_{1}(x)=\frac{\psi_{1}(x)}{\left\|\psi_{1}\right\|} \quad \text { where } \quad \psi_{1}(x)=g_{1}(x)
$$

and for $2 \leq r \leq n$ define

$$
\phi_{r}(x)=\frac{\psi_{r}(x)}{\left\|\psi_{r}\right\|} \quad \text { where } \quad \psi_{r}(x)=g_{r}(x)-\sum_{j=1}^{r-1} a_{r, j} \phi_{j}(x)
$$

and

$$
a_{r, j}=\left\langle g_{r}, \phi_{j}\right\rangle .
$$

Then, $\phi_{1}, \ldots, \phi_{n}$ form an orthonormal basis with respect to $P_{n}$.
Now, as before, we define

$$
\begin{equation*}
Z_{j}=\frac{1}{n} \sum_{i=1}^{n} Y_{i} \phi_{j}\left(x_{i}\right), \quad j=1, \ldots, n . \tag{8.27}
\end{equation*}
$$

It follows that

$$
Z_{j} \approx N\left(\theta_{j}, \frac{\sigma^{2}}{n}\right)
$$

and we can then use the methods that we developed in this chapter.

### 8.4 Density Estimation

Orthogonal function methods can also be used for density estimation. Let $X_{1}, \ldots, X_{n}$ be an IID sample from a distribution $F$ with density $f$ with support on $(0,1)$. We assume that $f \in L_{2}(0,1)$ so we can expand $f$ as

$$
f(x)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(x)
$$

where, as before, $\phi_{1}, \phi_{2}, \ldots$, is an orthogonal basis. Let

$$
\begin{equation*}
Z_{j}=\frac{1}{n} \sum_{i=1}^{n} \phi_{j}\left(X_{i}\right), \quad j=1,2, \ldots, n . \tag{8.28}
\end{equation*}
$$

Then,

$$
\mathbb{E}\left(Z_{j}\right)=\int \phi_{j}(x) f(x) d x=\theta_{j}
$$

and

$$
\mathbb{V}\left(Z_{j}\right)=\frac{1}{n}\left(\int \phi_{j}^{2}(x) f(x) d x-\theta_{j}^{2}\right) \equiv \sigma_{j}^{2} .
$$

As in the regression case, we take $\widehat{\theta}_{j}=0$ for $j>n$ and we estimate $\theta=$ $\left(\theta_{1}, \ldots, \theta_{n}\right)$ using a modulation estimator $\hat{\theta}=b Z=\left(b_{1} Z_{1}, \ldots, b_{n} Z_{n}\right)$. The risk of this estimator is

$$
\begin{equation*}
R(b)=\sum_{j=1}^{n} b_{j}^{2} \sigma_{j}^{2}+\sum_{j=1}^{n}\left(1-b_{j}\right)^{2} \theta_{j}^{2} . \tag{8.29}
\end{equation*}
$$

We estimate $\sigma_{j}^{2}$ by

$$
\widehat{\sigma}_{j}^{2}=\frac{1}{n^{2}} \sum_{i=1}^{n}\left(\phi_{j}\left(X_{i}\right)-Z_{j}\right)^{2}
$$

and $\theta_{j}^{2}$ by $Z_{j}^{2}-\widehat{\sigma}_{j}^{2}$; then we can estimate the risk by

$$
\begin{equation*}
\widehat{R}(b)=\sum_{j=1}^{n} b_{j}^{2} \widehat{\sigma}_{j}^{2}+\sum_{j=1}^{n}\left(1-b_{j}\right)^{2}\left(Z_{j}^{2}-\widehat{\sigma}_{j}^{2}\right)_{+} . \tag{8.30}
\end{equation*}
$$

Finally, we choose $\widehat{b}$ by minimizing $\widehat{R}(b)$ over some class of modulators $\mathcal{M}$. The density estimate can be negative. We can fix this by performing surgery: remove the negative part of the density and renormalize it to integrate to 1 . Better surgery methods are discussed in Glad et al. (2003).

### 8.5 Comparison of Methods

The methods we have introduced for nonparametric regression so far are local regression (Section 5.4), spline smoothing (Section 5.5), and orthogonal function smoothing. In many ways, these methods are very similar. They all involve a bias-variance tradeoff and they all require choosing a smoothing parameter. Local polynomial smoothers have the advantage that they automatically correct for boundary bias. It is possible to modify orthogonal function estimators to alleviate boundary bias by changing the basis slightly; see Efromovich (1999). An advantage of orthogonal function smoothing is that it converts nonparametric regression into the many Normal means problem, which is simpler, at least for theoretical purposes. There are rarely huge differences between the approaches, especially when these differences are judged relative to the width of confidence bands. Each approach has its champions and its detractors. It is wise to use all available tools for each problem. If they agree then the choice of method is one of convenience or taste; if they disagree then there is value in figuring out why they differ.

Finally, let us mention that there is a formal relationship between these approaches. For example, orthogonal function can be viewed as kernel smoothing with a particular kernel and vice versa. See Härdle et al. (1998) for details.

### 8.6 Tensor Product Models

The methods in this chapter extend readily to higher dimensions although our previous remarks about the curse of dimensionality apply here as well.

Suppose that $r\left(x_{1}, x_{2}\right)$ is a function of two variables. For simplicity, assume that $0 \leq x_{1}, x_{2} \leq 1$. If $\phi_{0}, \phi_{1}, \ldots$ is an orthonormal basis for $L_{2}(0,1)$ then the functions

$$
\left\{\phi_{j, k}\left(x_{1}, x_{2}\right)=\phi_{j}\left(x_{1}\right) \phi_{k}\left(x_{2}\right): j, k=0,1, \ldots,\right\}
$$

form an orthonormal basis for $L_{2}([0,1] \times[0,1])$, called the tensor product basis. The basis can be extended to $d$ dimensions in the obvious way.

Suppose that $\phi_{0}=1$. Then a function $r \in L_{2}([0,1] \times[0,1])$ can be expanded in the tensor product basis as

$$
\begin{aligned}
r\left(x_{1}, x_{2}\right) & =\sum_{j, k=0}^{\infty} \beta_{j, k} \phi_{j}\left(x_{1}\right) \phi_{k}\left(x_{2}\right) \\
& =\beta_{0}+\sum_{j=1}^{\infty} \beta_{j, 0} \phi_{j}\left(x_{1}\right)+\sum_{j=1}^{\infty} \beta_{0, j} \phi_{j}\left(x_{2}\right)+\sum_{j, k=1}^{\infty} \beta_{j, k} \phi_{j}\left(x_{1}\right) \phi_{k}\left(x_{2}\right) .
\end{aligned}
$$

This expansion has an ANOVA-like structure consisting of a mean, main effects, and interactions. This structure suggests a way to get better estimators. We could put stronger smoothness assumptions on higher-order interactions to get better rates of convergence (at the expense of more assumptions). See Lin (2000), Wahba et al. (1995), Gao et al. (2001), and Lin et al. (2000).

### 8.7 Bibliographic Remarks

The REACT method is developed in Beran (2000) and Beran and Dümbgen (1998). A different approach to using orthogonal functions is discussed in Efromovich (1999). REACT confidence sets are extended to nonconstant variance in Genovese et al. (2004), to wavelets in Genovese and Wasserman (2005), and to density estimation in Jang et al. (2004).

### 8.8 Exercises

1. Prove Lemma 8.4.
2. Prove Theorem 8.13.
3. Prove equation (8.19).
4. Prove equation (8.26).
5. Show that the estimator (8.12) is consistent.
6. Show that the estimator (8.12) is uniformly consistent over Sobolev ellipsoids.
7. Get the data on fragments of glass collected in forensic work from the book website. Let $Y$ be refractive index and let $x$ be aluminium content (the fourth variable). Perform a nonparametric regression to fit the model $Y=r(x)+\epsilon$. Use REACT and compare to local linear smoothing. Estimate the variance. Construct 95 percent confidence bands for your estimate.
8. Get the motorcycle data from the book website. The covariate is time (in milliseconds) and the response is acceleration at time of impact. Use REACT to fit the data. Compute 95 percent confidence bands. Compute a 95 percent confidence ball. Can you think of a creative way to display the confidence ball?
9. Generate 1000 observations from the model $Y_{i}=r\left(x_{i}\right)+\sigma \epsilon_{i}$ where $x_{i}=i / n, \epsilon_{i} \sim N(0,1)$ and $r$ is the Doppler function. Make three data sets corresponding to $\sigma=.1, \sigma=1$ and $\sigma=3$. Estimate the function using local linear regression and using REACT. In each case, compute a 95 percent confidence band. Compare the fits and the confidence bands.
10. Repeat the previous exercise but use Cauchy errors instead of Normal errors. How might you change the procedure to make the estimators more robust?
11. Generate $n=1000$ data points from $(1 / 2) N(0,1)+(1 / 2) N(\mu, 1)$. Compare kernel density estimators and REACT density estimators. Try $\mu=$ $0,1,2,3,4,5$.
12. Recall that a modulator is any vector of the form $b=\left(b_{1}, \ldots, b_{n}\right)$ such that $0 \leq b_{j} \leq 1, j=1, \ldots, n$. The greedy modulator is the modulator $b^{*}=\left(b_{1}^{*}, \ldots, b_{n}^{*}\right)$ chosen to minimize the risk $R(b)$ over all modulators.
(a) Find $b^{*}$.
(b) What happens if we try to estimate $b^{*}$ from the data? In particular, consider taking $\widehat{b}^{*}$ to minimize the estimated risk $\widehat{R}$. Why will this not work well? (The problem is that we are now minimizing $\widehat{R}$ over a very large class and $\widehat{R}$ does not approximate $R$ uniformly over such a large class.)
13. Let

$$
Y_{i}=r\left(x_{1 i}, x_{2 i}\right)+\epsilon_{i}
$$

where $\epsilon_{i} \sim N(0,1), x_{1 i}=x_{2 i}=i / n$ and $r\left(x_{1}, x_{2}\right)=x_{1}+\cos \left(x_{2}\right)$. Generate 1000 observations. Fit a tensor product model with $J_{1}$ basis elements for $x_{1}$ and $J_{2}$ basis elements for $x_{2}$. Use sure (Stein's unbiased risk estimator) to choose $J_{1}$ and $J_{2}$.
14. Download the air quality data set from the book website. Model ozone as a function of solar R , wind and temperature. Use a tensor product basis.

## 9

## Wavelets and Other Adaptive Methods

This chapter concerns estimating functions that are spatially inhomogeneous, functions $r(x)$ whose smoothness varies substantially with $x$. For example, Figure 9.1 shows the "blocks" function whose definition is given in Example 9.39. The function is very smooth except for several abrupt jumps. The top right plot shows 100 data points drawn according to $Y_{i}=r\left(x_{i}\right)+\epsilon_{i}$ with $\epsilon_{i} \sim N(0,1)$, and the $x_{i}$ s equally spaced.

It might be difficult to estimate $r$ using the methods we have discussed so far. If we use local regression with a large bandwidth, then we will smooth out the jumps; if we use a small bandwidth, then we will find the jumps but we will make the rest of the curve very wiggly. If we use orthogonal functions and keep only low-order terms, we will miss the jumps; if we allow higher-order terms we will find the jumps but we will make the rest of the curve very wiggly. The function estimates in Figure 9.1 illustrate this point. Another example of an inhomogeneous function is the Doppler function in Example 5.63.

Estimators that are designed to estimate such functions are said to be spatially adaptive or locally adaptive. A closely related idea is to find globally adaptive estimators; these are function estimators that perform well over large classes of function spaces. In this chapter we explore adaptive estimators with an emphasis on wavelet methods. In Section 9.9 we briefly consider some other adaptive methods.


FIGURE 9.1. The blocks function (upper left plot) is inhomogeneous. One-hundred data points are shown in the upper right plot. A local linear smoother with a small bandwidth (lower left) picks up the jumps but adds many wiggles. A local linear smoother with a large bandwidth (lower right) is smooth but misses the jumps.

Caveat Emptor! Before proceeding, a warning is in order. Adaptive estimation is difficult. Unless the signal-to-noise ratio is large, we cannot expect to adapt very well. To quote from Loader (1999b):

Locally adaptive procedures work well on examples with plenty of data, obvious structure and low noise. But . . . these are not difficult problems .... The real challenges ... occur when the structure is not obvious, and there are questions as to which features in a dataset are real. In such cases, simpler methods ... are most useful, and locally adaptive methods produce little benefit.

The reader is urged to keep this in mind as we proceed. See Section 9.10 for more discussion on this point. Despite the caveat, the methods in this chapter are important because they do work well in high signal-to-noise problems and, more importantly, the conceptual ideas underlying the methods are important in their own right.

The wavelet methods discussed in this chapter illustrate a concept that is becoming increasingly important in statistics and machine learning, namely, the notion of sparseness. A function $f=\sum_{j} \beta_{j} \phi_{j}$ is sparse in a basis $\phi_{1}, \phi_{2}, \ldots$ if most of the $\beta_{j}$ 's are zero (or close to zero). We will see that even some fairly complicated functions are sparse when expanded in a wavelet basis. Sparseness generalizes smoothness: smooth functions are sparse but there are also nonsmooth functions that are sparse. It is interesting to note that sparseness is not captured well by the $L_{2}$ norm but it is captured by the $L_{1}$ norm. For example, consider the $n$-vectors $a=(1,0, \ldots, 0)$ and $b=$ $(1 / \sqrt{n}, \ldots, 1 / \sqrt{n})$. Then both have the same $L_{2}$ norm: $\|a\|_{2}=\|b\|_{2}=1$. However, the $L_{1}$ norms are $\|a\|_{1}=\sum_{i}\left|a_{i}\right|=1$ and $\|b\|_{1}=\sum_{i}\left|b_{i}\right|=\sqrt{n}$. The $L_{1}$ norms reflects the sparsity of $a$.

Notation. Throughout this chapter, $\mathbb{Z}$ denotes the set of integers and $\mathbb{Z}_{+}$ denotes the set of positive integers. The Fourier transform $f^{*}$ of a function $f$ is

$$
\begin{equation*}
f^{*}(t)=\int_{-\infty}^{\infty} e^{-i x t} f(x) d x \tag{9.1}
\end{equation*}
$$

where $i=\sqrt{-1}$. If $f^{*}$ is absolutely integrable then $f$ can be recovered at almost all $x$ by the inverse Fourier transform

$$
\begin{equation*}
f(x)=\frac{1}{2 \pi} \int_{-\infty}^{\infty} e^{i x t} f^{*}(t) d t \tag{9.2}
\end{equation*}
$$

Given any function $f$ and integers $j$ and $k$, define

$$
\begin{equation*}
f_{j k}(x)=2^{j / 2} f\left(2^{j} x-k\right) \tag{9.3}
\end{equation*}
$$

### 9.1 Haar Wavelets

We begin with a simple wavelet called the Haar wavelet. The Haar father wavelet or Haar scaling function is defined by

$$
\phi(x)= \begin{cases}1 & \text { if } 0 \leq x<1  \tag{9.4}\\ 0 & \text { otherwise }\end{cases}
$$

The mother Haar wavelet is defined by

$$
\psi(x)=\left\{\begin{align*}
-1 & \text { if } 0 \leq x \leq \frac{1}{2}  \tag{9.5}\\
1 & \text { if } \frac{1}{2}<x \leq 1
\end{align*}\right.
$$



FIGURE 9.2. Some Haar wavelets. Left: the mother wavelet $\psi(x)$; right: $\psi_{2,2}(x)$.
For any integers $j$ and $k$ define $\phi_{j k}(x)$ and $\psi_{j k}(x)$ as in (9.3). The function $\psi_{j k}$ has the same shape as $\psi$ but it has been rescaled by a factor of $2^{j / 2}$ and shifted by a factor of $k$. See Figure 9.2 for some examples of Haar wavelets. Let

$$
W_{j}=\left\{\psi_{j k}, k=0,1, \ldots, 2^{j}-1\right\}
$$

be the set of rescaled and shifted mother wavelets at resolution $j$.
9.6 Theorem. The set of functions

$$
\left\{\phi, W_{0}, W_{1}, W_{2}, \ldots,\right\}
$$

is an orthonormal basis for $L_{2}(0,1)$.
It follows from this theorem that we can expand any function $f \in L_{2}(0,1)$ in this basis. Because each $W_{j}$ is itself a set of functions, we write the expansion as a double sum:

$$
\begin{equation*}
f(x)=\alpha \phi(x)+\sum_{j=0}^{\infty} \sum_{k=0}^{2^{j}-1} \beta_{j k} \psi_{j k}(x) \tag{9.7}
\end{equation*}
$$

where

$$
\alpha=\int_{0}^{1} f(x) \phi(x) d x, \quad \beta_{j k}=\int_{0}^{1} f(x) \psi_{j k}(x) d x
$$

We call $\alpha$ the scaling coefficient and the $\beta_{j k}$ 's are called the detail coefficients. We call the finite sum

$$
\begin{equation*}
f_{J}(x)=\alpha \phi(x)+\sum_{j=0}^{J-1} \sum_{k=0}^{2^{j}-1} \beta_{j k} \psi_{j k}(x) \tag{9.8}
\end{equation*}
$$

the resolution $J$ approximation to $f$. The total number of terms in this sum is

$$
1+\sum_{j=0}^{J-1} 2^{j}=1+2^{J}-1=2^{J}
$$

9.9 Example. Figure 9.3 shows the Doppler signal (Example 5.63) and its resolution $J$ approximation for $J=3,5$ and $J=8$.


FIGURE 9.3. The Doppler signal (top left) and its reconstruction $f_{J}(x)=\alpha \phi(x)+\sum_{j=0}^{J-1} \sum_{k} \beta_{j k} \psi_{j k}(x)$ based on $J=3$ (top right), $J=5$ (bottom left) and $J=8$ (bottom right).

When $j$ is large, $\psi_{j k}$ is a very localized function. This makes it possible to add a blip to a function in one place without adding wiggles elsewhere. This is what makes a wavelet basis a good tool for modeling inhomogeneous functions.

Figure 9.4 shows that blocks function and the coefficients of the expansion of the function in the Haar basis. Notice that the expansion is sparse (most coefficients are zero) since nonzero coefficients are needed mainly where the function has jumps.

mother coefficients (level 2)

mother coefficients (level 3)

FIGURE 9.4. The blocks function $f(x)$ is shown on top. The second plot show the coefficients of the father wavelets. The third plot show the coefficients of the sec-ond-level mother wavelets. The remaining plots show the coefficients of the mother wavelets at higher levels. Despite the fact that the function is not smooth, the function is sparse: most coefficients are zero.

### 9.2 Constructing Wavelets

Haar wavelets are useful because they are localized, that is, they have bounded support. But Haar wavelets are not smooth. How can we construct other wavelets? In particular, how can we construct localized, smooth wavelets? The answer is quite involved. We will give a brief synopsis of the main ideas; see Härdle et al. (1998) and Daubechies (1992) for more details.

Given any function $\phi$ define

$$
\begin{align*}
V_{0} & =\left\{f: f(x)=\sum_{k \in \mathbb{Z}} c_{k} \phi(x-k), \sum_{k \in \mathbb{Z}} c_{k}^{2}<\infty\right\} \\
V_{1} & =\left\{f(x)=g(2 x): g \in V_{0}\right\}  \tag{9.10}\\
V_{2} & =\left\{f(x)=g(2 x): g \in V_{1}\right\} \\
\vdots & \vdots
\end{align*}
$$

9.11 Definition. Given a function $\phi$, define $V_{0}, V_{1}, \ldots$, as in (9.10). We say that $\phi$ generates a multiresolution analysis (MRA) of $\mathbb{R}$ if

$$
\begin{equation*}
V_{j} \subset V_{j+1}, \quad j \geq 0 \tag{9.12}
\end{equation*}
$$

and

$$
\begin{equation*}
\bigcup_{j \geq 0} V_{j} \quad \text { is dense in } L_{2}(\mathbb{R}) \tag{9.13}
\end{equation*}
$$

We call $\phi$ the father wavelet or the scaling function.
Equation (9.13) means that for any function $f \in L_{2}(\mathbb{R})$, there exists a sequence of functions $f_{1}, f_{2}, \ldots$ such that each $f_{r} \in \bigcup_{j \geq 0} V_{j}$ and $\left\|f_{r}-f\right\| \rightarrow 0$ as $r \rightarrow \infty$.
9.14 Lemma. If $V_{0}, V_{1}, \ldots$ is an MRA generated by $\phi$, then $\left\{\phi_{j k}, k \in \mathbb{Z}\right\}$ is an orthonormal basis for $V_{j}$.
9.15 Example. If $\phi$ is the father Haar wavelet, then $V_{j}$ is the set of functions $f \in L_{2}(\mathbb{R})$ that are piecewise constant on $\left[k 2^{-j},(k+1) 2^{-j}\right)$ for $k \in \mathbb{Z}$. It is easy to check that $V_{0}, V_{1}, \ldots$ form an MRA.

Suppose we have an MRA. Since $\phi \in V_{0}$ and $V_{0} \subset V_{1}$, we also have that $\phi \in V_{1}$. Since $\left\{\phi_{1 k}, k \in \mathbb{Z}\right\}$ is an orthonormal basis for $V_{1}$, we can thus write
$\phi$ as a linear combination of functions in $V_{1}$ :

$$
\begin{equation*}
\phi(x)=\sum_{k} \ell_{k} \phi_{1 k}(x) \tag{9.16}
\end{equation*}
$$

where $\ell_{k}=\int \phi(x) \phi_{1 k}(x) d x$ and $\sum_{k} \ell_{k}^{2}<\infty$. Equation (9.16) is called the two-scale relationship or the dilation equation. For Haar wavelets $\ell_{0}=$ $\ell_{1}=2^{-1 / 2}$ and $\ell_{k}=0$ for $k \neq 0,1$. The coefficients $\left\{\ell_{k}\right\}$ are called scaling coefficients. The two-scale relationship implies that

$$
\begin{equation*}
\phi^{*}(t)=m_{0}(t / 2) \phi^{*}(t / 2) \tag{9.17}
\end{equation*}
$$

where $m_{0}(t)=\sum_{k} \ell_{k} e^{-i t k} / \sqrt{2}$. Applying the above formula recursively, we see that $\phi^{*}(t)=m_{0}(t / 2) \prod_{k=1}^{\infty} m_{0}\left(t / 2^{k}\right) \phi^{*}(0)$. This suggests that, given just the scaling coefficients, we can compute $\phi^{*}(t)$ and then take the inverse Fourier transform to find $\phi(t)$. An example of how to construct a father wavelet from a set of scaling coefficients is given in the next theorem.
9.18 Theorem. Given coefficients $\left\{\ell_{k}, k \in \mathbb{Z}\right\}$, define a function

$$
\begin{equation*}
m_{0}(t)=\frac{1}{\sqrt{2}} \sum_{k} \ell_{k} e^{-i t k} \tag{9.19}
\end{equation*}
$$

Let

$$
\begin{equation*}
\phi^{*}(t)=\prod_{j=1}^{\infty} m_{0}\left(\frac{t}{2^{j}}\right) \tag{9.20}
\end{equation*}
$$

and let $\phi$ be the inverse Fourier transform of $\phi^{*}$. Suppose that

$$
\begin{equation*}
\frac{1}{\sqrt{2}} \sum_{k=N_{0}}^{N_{1}} \ell_{k}=1 \tag{9.21}
\end{equation*}
$$

for some $N_{0}<N_{1}$, that

$$
\left|m_{0}(t)\right|^{2}+\left|m_{0}(t+\pi)\right|^{2}=1
$$

that $m_{0}(t) \neq 0$ for $|t| \leq \pi / 2$, and that there exists a bounded nonincreasing function $\Phi$ such that $\int \Phi(|u|) d u<\infty$ and

$$
|\phi(x)| \leq \Phi(|x|)
$$

for almost all $x$. Then $\phi$ is a compactly supported father wavelet and $\phi$ is zero outside the interval $\left[N_{0}, N_{1}\right]$.

Next, define $W_{k}$ to be the orthogonal complement of $V_{k}$ in $V_{k+1}$. In other words, each $f \in V_{k+1}$ can be written as a sum $f=v_{k}+w_{k}$ where $v_{k} \in V_{k}$, $w_{k} \in W_{k}$, and $v_{k}$ and $w_{k}$ are orthogonal. We write

$$
V_{k+1}=V_{k} \bigoplus W_{k}
$$

Thus,

$$
L_{2}(\mathbb{R})=\overline{\bigcup_{k} V_{k}}=V_{0} \bigoplus W_{0} \bigoplus W_{1} \bigoplus \cdots
$$

Define the mother wavelet by

$$
\psi(x)=\sqrt{2} \sum_{k}(-1)^{k+1} \ell_{1-k} \phi(2 x-k) .
$$

9.22 Theorem. The functions $\left\{\psi_{j k}, k \in \mathbb{Z}\right\}$ form a basis for $W_{j}$. The functions

$$
\begin{equation*}
\left\{\phi_{k}, \psi_{j k}, k \in \mathbb{Z}, j \in \mathbb{Z}_{+}\right\} \tag{9.23}
\end{equation*}
$$

are an orthonormal basis for $L_{2}(\mathbb{R})$. Hence, any $f \in L_{2}$ can be written as

$$
\begin{equation*}
f(x)=\sum_{k} \alpha_{0 k} \phi_{0 k}(x)+\sum_{j=0}^{\infty} \sum_{k} \beta_{j k} \psi_{j k}(x) \tag{9.24}
\end{equation*}
$$

where

$$
\alpha_{0 k}=\int f(x) \phi_{0 k}(x) d x \quad \text { and } \quad \beta_{j k}=\int f(x) \psi_{j k}(x) d x
$$

We have been denoting the first space in the MRA by $V_{0}$. This is simply a convention. We could just as well denote it by $V_{j_{0}}$ for some integer $j_{0}$ in which case we write

$$
f(x)=\sum_{k} \alpha_{j_{0} k} \phi_{j_{0} k}(x)+\sum_{j=j_{0}}^{\infty} \sum_{k} \beta_{j k} \psi_{j k}(x) .
$$

Of course, we still have not explained how to choose scaling coefficients to create a useful wavelet basis. We will not discuss the details here, but clever choices of scaling coefficients lead to wavelets with desirable properties. For example, in 1992 Ingrid Daubechies constructed a smooth, compactly supported "nearly" symmetric ${ }^{1}$ wavelet called a symmlet. Actually, this is a

[^20]family of wavelets. A father symmlet of order $N$ has support [ $0,2 N-1$ ] while the mother has support $[-N+1, N]$. The mother has $N$ vanishing moments (starting with the $0^{\text {th }}$ moment). The higher $N$ is, the smoother the wavelet. There is no closed form expression for this wavelet (or for most wavelets) but it can be computed rapidly. Figure 9.5 shows the symmlet 8 mother wavelet which we will use in our examples. The scaling coefficients are:
\[

$$
\begin{array}{rrrr}
0.0018899503 & -0.0003029205 & -0.0149522583 & 0.0038087520 \\
0.0491371797 & -0.0272190299 & -0.0519458381 & 0.3644418948 \\
0.7771857517 & 0.4813596513 & -0.0612733591 & -0.1432942384 \\
0.0076074873 & 0.0316950878 & -0.0005421323 & -0.0033824160
\end{array}
$$
\]



FIGURE 9.5. The symmlet 8 mother wavelet.

### 9.3 Wavelet Regression

Consider the regression problem

$$
\begin{equation*}
Y_{i}=r\left(x_{i}\right)+\sigma \epsilon_{i} \tag{9.25}
\end{equation*}
$$

where $\epsilon_{i} \sim N(0,1)$ and $x_{i}=i / n$. We will further assume that $n=2^{J}$ for some $J$. We will need to do some boundary corrections for $x$ near 0 or 1 ; we defer that issue until later.

To estimate $r$ using wavelets, we proceed as follows: First, approximate $r$ with an expansion using $n$ terms:

$$
\begin{equation*}
r(x) \approx r_{n}(x)=\sum_{k=0}^{2^{j_{0}}-1} \alpha_{j_{0}, k} \phi_{j_{0}, k}(x)+\sum_{j=j_{0}}^{J} \sum_{k=0}^{2^{j}-1} \beta_{j k} \psi_{j k}(x) \tag{9.26}
\end{equation*}
$$

where $\alpha_{j_{0}, k}=\int r(x) \phi_{j_{0}, k}(x) d x$ and $\beta_{j k}=\int r(x) \psi_{j k}(x) d x$. We refer to $\left\{\beta_{j k}\right.$, $k=0,1, \ldots\}$ as the level-j coefficients. Form preliminary estimates of the coefficients: ${ }^{2}$

$$
\begin{equation*}
S_{k}=\frac{1}{n} \sum_{i} \phi_{j_{0}, k}\left(x_{i}\right) Y_{i} \quad \text { and } \quad D_{j k}=\frac{1}{n} \sum_{i} \psi_{j k}\left(x_{i}\right) Y_{i} \tag{9.27}
\end{equation*}
$$

which are called the empirical scaling coefficients and the empirical detail coefficients. Arguing as in the last chapter, we have that

$$
S_{k} \approx N\left(\alpha_{j_{0}, k}, \frac{\sigma^{2}}{n}\right) \quad \text { and } \quad D_{j k} \approx N\left(\beta_{j k}, \frac{\sigma^{2}}{n}\right)
$$

Next we use the following robust estimator for $\sigma$ :

$$
\widehat{\sigma}=\sqrt{n} \times \frac{\operatorname{median}\left(\left|D_{J-1, k}-\operatorname{median}\left(D_{J-1, k}\right)\right|: k=0, \ldots, 2^{J-1}-1\right)}{0.6745}
$$

Inhomogeneous functions may have a few large wavelet coefficients even at the highest level $J$ and this robust estimator should be relatively insensitive to such coefficients.

For the scaling coefficients, we take

$$
\widehat{\alpha}_{j_{0}, k}=S_{k}
$$

[^21]To estimate the coefficients $\beta_{j k}$ of the mother wavelets we use a special type of shrinkage on the $D_{j k}$ called thresholding which we describe in more detail in the next section. Finally, we plug the estimates into (9.26):

$$
\widehat{r}_{n}(x)=\sum_{k=0}^{2^{j 0}-1} \widehat{\alpha}_{j_{0}, k} \phi_{j_{0}, k}(x)+\sum_{j=j_{0}}^{J} \sum_{k=0}^{2^{j}-1} \widehat{\beta}_{j k} \psi_{j k}(x)
$$

### 9.4 Wavelet Thresholding

The wavelet regression method is the same as the procedure we used in the last chapter except for two changes: the basis is different and we use a different type of shrinkage called thresholding in which $\widehat{\beta}_{j k}$ is set to 0 if $D_{j k}$ is small. Thresholding works better at finding jumps in the function than linear shrinkage. To see why, think of a function that is smooth except that it has some jumps in a few places. If we expand this function in a wavelet basis, the coefficients will be sparse, meaning that most will be small except for a few coefficients corresponding to the location of the jumps. Intuitively, this suggests setting most of the estimated coefficients to zero except the very large ones. This is precisely what a threshold rule does. More formally, we will see in the next section that thresholding shrinkage yields minimax estimators over large function spaces. Here are the details.

The estimates of the coefficients $\alpha_{j_{0}, k}$ of the father wavelet are equal to the empirical coefficients $S_{k}$; no shrinkage is applied. The estimates of the coefficients of the mother wavelets are based on shrinking the $D_{j k}$ 's as follows. Recall that

$$
\begin{equation*}
D_{j k} \approx \beta_{j k}+\frac{\sigma}{\sqrt{n}} \epsilon_{j k} \tag{9.28}
\end{equation*}
$$

The linear shrinkers we used in Chapters 7 and 8 were of the form $\widehat{\beta}_{j k}=c D_{j k}$ for some $0 \leq c \leq 1$. For wavelets, we use nonlinear shrinkage called thresholding which comes in two flavors: hard thresholding and soft thresholding. The hard threshold estimator is

$$
\widehat{\beta}_{j k}= \begin{cases}0 & \text { if }\left|D_{j k}\right|<\lambda  \tag{9.29}\\ D_{j k} & \text { if }\left|D_{j k}\right| \geq \lambda\end{cases}
$$

The soft threshold estimator is

$$
\widehat{\beta}_{j k}= \begin{cases}D_{j k}+\lambda & \text { if } D_{j k}<-\lambda  \tag{9.30}\\ 0 & \text { if }-\lambda \leq D_{j k}<\lambda \\ D_{j k}-\lambda & \text { if } D_{j k}>\lambda\end{cases}
$$



Hard thresholding


Soft thresholding

FIGURE 9.6. Hard and soft thresholding.
which can be written more succinctly as

$$
\begin{equation*}
\widehat{\beta}_{j k}=\operatorname{sign}\left(D_{j k}\right)\left(\left|D_{j k}\right|-\lambda\right)_{+} . \tag{9.31}
\end{equation*}
$$

See Figure 9.6. In either case, the effect is to keep large coefficients and set the others equal to 0 .

We will restrict attention to soft thresholding. We still need to choose the threshold $\lambda$. There are several methods for choosing $\lambda$. The simplest rule is the universal threshold defined by

$$
\begin{equation*}
\lambda=\widehat{\sigma} \sqrt{\frac{2 \log n}{n}} \tag{9.32}
\end{equation*}
$$

To understand the intuition behind this rule, consider what happens when there is no signal, that is, when $\beta_{j k}=0$ for all $j$ and $k$. In this case, we would like all $\widehat{\beta}_{j k}$ to be 0 with high probability.
9.33 Theorem. Suppose that $\beta_{j k}=0$ for all $j$ and $k$ and let $\widehat{\beta}_{j k}$ be the softthreshold estimator with the universal threshold (9.32). Then

$$
\mathbb{P}\left(\widehat{\beta}_{j k}=0 \text { for all } j, k\right) \rightarrow 1
$$

as $n \rightarrow \infty$.

Proof. To simplify the proof, assume that $\sigma$ is known. Now, $D_{j k} \approx$ $N\left(0, \sigma^{2} / n\right)$. Recall Mills' inequality: if $Z \sim N(0,1)$ then $\mathbb{P}(|Z|>t) \leq(c / t)$ $e^{-t^{2} / 2}$ where $c=\sqrt{2 / \pi}$ is a constant. Thus,

$$
\begin{aligned}
\mathbb{P}\left(\max \left|D_{j k}\right|>\lambda\right) & \leq \sum_{j, k} \mathbb{P}\left(\left|D_{j k}\right|>\lambda\right) \leq \sum_{j, k} \mathbb{P}\left(\frac{\sqrt{n}\left|D_{j k}\right|}{\sigma}>\frac{\sqrt{n} \lambda}{\sigma}\right) \\
& =\sum_{j, k} \frac{c \sigma}{\lambda \sqrt{n}} \exp \left\{-\frac{1}{2} \frac{n \lambda^{2}}{\sigma^{2}}\right\}=\frac{c}{\sqrt{2 \log n}} \rightarrow 0 .
\end{aligned}
$$

More support for the thresholding rule is given by the following theorem which shows that soft thresholding does nearly as well as an "oracle" that chooses the best threshold.
9.34 Theorem (Donoho and Johnstone 1994). Let

$$
Y_{i}=\theta_{i}+\frac{\sigma}{\sqrt{n}} \epsilon_{i}, \quad i=1, \ldots, n
$$

where $\epsilon_{i} \sim N(0,1)$. For each $S \subset\{1, \ldots, n\}$ define the kill it or keep it estimator

$$
\widehat{\theta}_{S}=\left(X_{1} I(1 \in S), \ldots, X_{n} I(n \in S)\right)
$$

Define the oracle risk

$$
\begin{equation*}
R_{n}^{*}=\min _{S} R\left(\widehat{\theta}_{S}, \theta\right) \tag{9.35}
\end{equation*}
$$

where the minimum is over all kill it or keep it estimators, that is, $S$ varies over all subsets of $\{1, \ldots, n\}$. Then

$$
\begin{equation*}
R_{n}^{*}=\sum_{i=1}^{n}\left(\theta_{i}^{2} \wedge \frac{\sigma^{2}}{n}\right) \tag{9.36}
\end{equation*}
$$

Further, if

$$
\widehat{\theta}=\left(t\left(X_{1}\right), \ldots, t\left(X_{n}\right)\right)
$$

where $t(x)=\operatorname{sign}(x)\left(|x|-\lambda_{n}\right)_{+}$and $\lambda_{n}=\sigma \sqrt{2 \log n / n}$, then for every $\theta \in \mathbb{R}^{n}$,

$$
\begin{equation*}
R_{n}^{*} \leq R(\widehat{\theta}, \theta) \leq(2 \log n+1)\left(\frac{\sigma^{2}}{n}+R_{n}^{*}\right) \tag{9.37}
\end{equation*}
$$

Donoho and Johnstone call wavelet regression with the universal threshold, VisuShrink. Another estimator, called SureShrink is obtained by using a different threshold $\lambda_{j}$ for each level. The threshold $\lambda_{j}$ is chosen to minimize SURE (see Section 7.4) which, in this case, is

$$
\begin{equation*}
S\left(\lambda_{j}\right)=\sum_{k=1}^{n_{j}}\left[\frac{\widehat{\sigma}^{2}}{n}-2 \frac{\widehat{\sigma}^{2}}{n} I\left(\left|\widetilde{\beta}_{j k}\right| \leq \lambda_{j}\right)+\min \left(\widetilde{\beta}_{j k}^{2}, \lambda_{j}^{2}\right)\right] \tag{9.38}
\end{equation*}
$$

where $n_{j}=2^{j-1}$ is the number of parameters at level $j$. The minimization is performed over $0 \leq \lambda_{j} \leq\left(\hat{\sigma} / \sqrt{n_{j}}\right) \sqrt{2 \log n_{j}} .{ }^{3}$
9.39 Example. The "blocks" function-introduced by Donoho and Johnstone (1995) -is defined by $r(x)=3.655606 \times \sum_{j=1}^{11} h_{j} K\left(x-t_{j}\right)$ where

$$
\begin{aligned}
t & =(0.10,0.13,0.15,0.23,0.25,0.40,0.44,0.65,0.76,0.78,0.81) \\
h & =(4,-5,3,-4,5,-4.2,2.1,4.3,-3.1,2.1,-4.2)
\end{aligned}
$$

The top left panel in Figure 9.7 shows $r(x)$. The top right plot shows 2048 data points generates from $Y_{i}=r(i / n)+\epsilon_{i}$ with $\epsilon \sim N(0,1)$. The bottom left plot is the wavelet estimator using soft-thresholding with a universal threshold. We used a symmlet 8 wavelet. The bottom right plot shows a local linear fit with bandwidth chosen by cross-validation. The wavelet estimator is slightly better since the local linear fit has some extra wiggliness. However, the difference is not dramatic. There are situations where even this small difference might matter, such as in signal processing. But for ordinary nonparametric regression problems, there is not much practical difference between the estimators. Indeed, if we added confidence bands to these plots, they would surely be much wider than the difference between these estimators.

This is an easy example in the sense that the basic shape of the curve is evident from a plot of the data because the noise level is low. Let us now consider a noisier version of this example. We increase $\sigma$ to 3 and we reduce the sample size to 256 . The results are shown in Figure 9.8. It is hard to say which estimator is doing better here. Neither does particularly well.

### 9.5 Besov Spaces

Wavelet threshold regression estimators have good optimality properties over
Besov spaces which we now define. Let

$$
\Delta_{h}^{(r)} f(x)=\sum_{k=0}^{r}\binom{r}{k}(-1)^{k} f(x+k h)
$$

Thus, $\Delta_{h}^{(0)} f(x)=f(x)$ and

$$
\Delta_{h}^{(r)} f(x)=\Delta_{h}^{(r-1)} f(x+h)-\Delta_{h}^{(r-1)} f(x)
$$

[^22]

FIGURE 9.7. The "blocks" function from Example 9.39. Top left: the function $r(x)$. Top right: 2048 data points. Bottom left: $\widehat{r}_{n}$ using wavelets. Bottom right: $\widehat{r}_{n}$ using local linear regression with bandwidth chosen by cross-validation.


FIGURE 9.8. The "blocks" function from Example 9.39. Top left: the function $r(x)$. Top right: 256 data points. Bottom left: $\widehat{r}_{n}$ using wavelets. Bottom right: $\widehat{r}_{n}$ using local linear regression with bandwidth chosen by cross-validation.

Next define

$$
w_{r, p}(f ; t)=\sup _{|h| \leq t}\left\|\Delta_{h}^{(r)} f\right\|_{p}
$$

where $\|g\|_{p}=\left\{\int|g(x)|^{p} d x\right\}^{1 / p}$. Given $(p, q, \varsigma)$, let $r$ be such that $r-1 \leq \varsigma \leq r$. The Besov seminorm is defined by

$$
|f|_{p, q}^{\varsigma}=\left[\int_{0}^{\infty}\left(h^{-\varsigma} w_{r, p}(f ; h)\right)^{q} \frac{d h}{h}\right]^{1 / q} .
$$

For $q=\infty$ we define

$$
|f|_{p, \infty}^{\varsigma}=\sup _{0<h<1} \frac{w_{r, p}(f ; h)}{h^{\varsigma}} .
$$

The Besov space $B_{p, q}^{\varsigma}(c)$ is defined to be the set of functions $f$ mapping $[0,1]$ into $\mathbb{R}$ such that $\int|f|^{p}<\infty$ and $|f|_{p, q}^{\varsigma} \leq c$.

This definition is hard to understand but here are some special cases. The Sobolev space $W(m)$ (see Definition 7.8) corresponds to the Besov ball $B_{2,2}^{m}$. The generalized Sobolev space $W_{p}(m)$ which uses an $L_{p}$ norm on the $m^{\text {th }}$ derivative is almost a Besov space in the sense that $B_{p, 1}^{m} \subset W_{p}(m) \subset B_{p, \infty}^{m}$. Let $s=m+\delta$ for some integer $m$ and some $\delta \in(0,1)$. The Hölder space is the set of bounded functions with bounded $m^{t h}$ derivative such that $\left|f^{m}(u)-f^{m}(t)\right| \leq$ $|u-t|^{\delta}$ for all $u, t$. This space is equivalent to $B_{\infty, \infty}^{m+\delta}$. The set $T$ consisting of functions of bounded variation satisfies: $B_{1,1}^{1} \subset T \subset B_{1, \infty}^{1}$. Thus, Besov spaces include a wide range of familiar function spaces.

It is easier to understand Besov spaces in terms of the coefficients of the wavelet expansion. If the wavelets are sufficiently smooth then the wavelet coefficients $\beta$ of a function $f \in B_{p, q}^{\varsigma}(c)$ satisfy $\|\beta\|_{p, q}^{\varsigma} \leq c$, where

$$
\begin{equation*}
\|\beta\|_{p, q}^{\varsigma}=\left(\sum_{j=j_{0}}^{\infty}\left(2^{j(\varsigma+(1 / 2)-(1 / p))}\left(\sum_{k}\left|\beta_{j k}\right|^{p}\right)^{1 / p}\right)^{q}\right)^{1 / q} \tag{9.40}
\end{equation*}
$$

In the theorem that follows, we use the following notation: write $a_{n} \asymp b_{n}$ to mean that $a_{n}$ and $b_{n}$ go to 0 at the same rate. Formally,

$$
0<\liminf _{n \rightarrow \infty}\left|\frac{a_{n}}{b_{n}}\right| \leq \limsup _{n \rightarrow \infty}\left|\frac{a_{n}}{b_{n}}\right|<\infty
$$

9.41 Theorem (Donoho-Johnstone 1995). Let $\widehat{r}_{n}$ be the SureShrink estimator. Let $\psi$ have $r$ null moments and $r$ continuous derivatives where $r>$ $\max \{1, \varsigma\}$. Let $R_{n}(p, q, \varsigma, C)$ denote the minimax risk over the Besov ball $B_{p, q}^{\varsigma}(c)$. Then

$$
\sup _{r \in B_{p, q}^{\varsigma}(C)} \frac{1}{n} \mathbb{E}\left(\sum_{i}\left(\widehat{r}_{n}\left(x_{i}\right)-r\left(x_{i}\right)\right)^{2}\right) \asymp R_{n}(p, q, \varsigma, c)
$$

for all $1 \leq p, q \leq \infty, C \in(0, \infty)$ and $\varsigma_{0}<\varsigma<r$ where

$$
\varsigma_{0}=\max \left\{\frac{1}{p}, 2\left(\frac{1}{p}-\frac{1}{2}\right)^{+}\right\} .
$$

No linear estimator achieves the optimal rate over this range of spaces. The universal shrinkage rule also attains the minimax rate, up to factors involving $\log n$.

The theorem says that wavelet estimators based on threshold rules achieve the minimax rate, simultaneously over a large set of Besov spaces. Other estimators, such as local regression estimators with a constant bandwidth, do not have this property.

### 9.6 Confidence Sets

At the time of this writing, practical, simultaneous confidence bands for wavelet estimators are not available. An asymptotic pointwise method is given in Picard and Tribouley (2000). Confidence balls are available from Genovese and Wasserman (2005).

For a Besov space $B_{p, q}^{\varsigma}$, let

$$
\gamma= \begin{cases}\varsigma & p \geq 2  \tag{9.42}\\ \varsigma+\frac{1}{2}-\frac{1}{p} & 1 \leq p<2\end{cases}
$$

Let $B(p, q, \gamma)$ be the set of wavelet coefficients corresponding to functions in the Besov space. We assume that the mother and father wavelets are bounded, have compact support, and have derivatives with finite $L_{2}$ norms. Let $\mu^{n}=$ $\left(\mu_{1}, \ldots, \mu_{n}\right)$ be the first $n$ wavelet coefficients strung out as a single vector.
9.43 Theorem (Genovese and Wasserman 2004). Let $\widehat{\mu}$ be the estimated wavelet coefficients using the universal soft threshold $\lambda=\widehat{\sigma} \sqrt{\log n / n}$. Define

$$
\begin{equation*}
D_{n}=\left\{\mu^{n}: \sum_{\ell=1}^{n}\left(\mu_{\ell}-\widehat{\mu}_{\ell}\right)^{2} d x \leq s_{n}^{2}\right\} \tag{9.44}
\end{equation*}
$$

where

$$
\begin{gather*}
s_{n}^{2}=\sqrt{2} \sigma^{2} \frac{z_{\alpha}}{\sqrt{n}}+S_{n}(\lambda)  \tag{9.45}\\
S_{n}(\lambda)=\frac{\widehat{\sigma}^{2}}{n} 2^{j_{0}}+\sum_{j=j_{0}}^{J} S\left(\lambda_{j}\right) \tag{9.46}
\end{gather*}
$$

and

$$
S_{j}\left(\lambda_{j}\right)=\sum_{k=1}^{n_{j}}\left[\frac{\sigma^{2}}{n}-2 \frac{\sigma^{2}}{n} I\left(\left|\widetilde{\beta}_{j k}\right| \leq \lambda_{j}\right)+\min \left(\widetilde{\beta}_{j k}^{2}, \lambda_{j}^{2}\right)\right]
$$

Then, for any $\delta>0$,

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \sup _{\mu \in \Delta(\delta)}\left|\mathbb{P}\left(\mu^{n} \in D_{n}\right)-(1-\alpha)\right|=0 \tag{9.47}
\end{equation*}
$$

where

$$
\Delta(\delta)=\bigcup\{B(p, q, \gamma): p \geq 1, q \geq 1, \gamma>(1 / 2)+\delta\}
$$

### 9.7 Boundary Corrections and Unequally Spaced Data

If the data live on an interval but the wavelet basis lives on $\mathbb{R}$, then a correction is needed since the wavelets will not usually be orthonormal when restricted to the interval. The simplest approach is to mirror the data. The data are repeated in reverse order around the endpoint. Then the previously discussed methods are applied.

When the data are not equally spaced or $n$ is not a power of 2 , we can put the data into equally spaced bins and average the data over the bins. We choose the bins as small as possible subject to having data in each bin and subject to the number of bins $m$ being of the form $m=2^{k}$ for some integer $k$.

For other approaches to these problems see Härdle et al. (1998) as well as the references therein.

### 9.8 Overcomplete Dictionaries

Although wavelet bases are very flexible, there are cases where one might want to build even richer bases. For example, one might consider combining several bases together. This leads to the idea of a dictionary.

Let $Y=\left(Y_{1}, \ldots, Y_{n}\right)^{T}$ be a vector of observations where $Y_{i}=r\left(x_{i}\right)+\epsilon_{i}$. Let $D$ be an $n \times m$ matrix with $m>n$. Consider estimating $r$ using $D \beta$ where $\beta=\left(\beta_{1}, \ldots, \beta_{m}\right)^{T}$. If $m$ were equal to $n$ and the columns of $D$ were orthonormal, we are back to orthogonal basis regression as in this and the last chapter. But when $m>n$, the columns can no longer be orthogonal and we say that the dictionary is overcomplete. For example, we might take $D$ to have $m=2 n$ columns: the first $n$ columns being the basis elements of a cosine
basis and the second $n$ columns being the basis elements of a "spike" basis. This would enable us to estimate a function that is "smooth plus spikes."

There is growing evidence - theoretical and practical-that the method called basis pursuit, due to Chen et al. (1998), leads to good estimators. In this method, one chooses $\widehat{\beta}$ to minimize

$$
\|Y-D \beta\|_{2}^{2}+\lambda\|\beta\|_{1}
$$

where $\|\cdot\|_{2}$ denotes the $L_{2}$ norm, $\|\cdot\|_{1}$ denotes the $L_{1}$ norm and $\lambda>0$ is a constant. Basis pursuit is related to the regression variable selection techniques called the lasso (Tibshirani (1996)) and LARS (Efron et al. (2004)).

### 9.9 Other Adaptive Methods

There are many other spatially adaptive methods besides wavelets; these include variable bandwidth kernel methods (Lepski et al. (1997), Müller and Stadtmller (1987)), local polynomials (Loader (1999a) and Fan and Gijbels (1996)) variable knot splines (Mammen and van de Geer (1997)) among others. Here we outline an especially simple, elegant method due to Goldenshluger and Nemirovski (1997) called intersecting confidence intervals (ICI).

Consider estimating a regression function $r(x)$ at a point $x$ and let $Y_{i}=$ $r\left(x_{i}\right)+\epsilon_{i}$ where $\epsilon_{i} \sim N\left(0, \sigma^{2}\right)$. To simplify the discussion, we will assume that $\sigma$ is known. The idea is to construct a confidence interval for $r(x)$ using a sequence of increasing bandwidths $h$. We will then choose the first bandwidth at which the intervals do not intersect. See Figure 9.9.

Let

$$
\widehat{r}_{h}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x, h)
$$

be a linear estimator depending on bandwidth $h$. Let the bandwidth $h$ vary in a finite set $h_{1}<\cdots<h_{n}$ and let $\widehat{r}_{j}=\widehat{r}_{h_{j}}(x)$. For example, with equally spaced data on $[0,1]$ we might take $h_{j}=j / n$. We can write

$$
\widehat{r}_{j}=\sum_{i=1}^{n} r\left(x_{i}\right) \ell_{i}\left(x, h_{j}\right)+\xi_{j}
$$

where

$$
\xi_{j}=\sum_{i=1}^{n} \epsilon_{i} \ell_{i}\left(x, h_{j}\right) \sim N\left(0, s_{j}^{2}\right), \quad s_{j}=\sigma \sqrt{\sum_{i=1}^{n} \ell_{i}^{2}\left(x, h_{j}\right)} .
$$



FIGURE 9.9. The Goldenshluger and Nemirovski ICI method. All the intervals $\left\{D_{j}: j \leq j^{*}\right\}$ contain the true value $r(x)$. All the intervals $\left\{D_{j}: j \leq j^{+}\right\}$overlap. The estimate $\widehat{r}(x)$ is no more than $2 \rho_{j^{+}}+4 \rho_{j^{*}} \leq 6 \rho_{j^{*}}$ away from $r(x)$.

Let $\rho_{j}=\kappa s_{j}$ where $\kappa>\sqrt{2 \log n}$. Then

$$
\mathbb{P}\left(\max _{j}\left|\xi_{j}\right|>\rho_{j}, \text { for some } j\right) \leq n P(|N(0,1)|>\kappa) \rightarrow 0 \text { as } n \rightarrow \infty
$$

So $\left|\xi_{j}\right| \leq \rho_{j}$ for all $j$ expect on a set of probability tending to 0 . For the rest of the argument, assume that $\left|\xi_{j}\right| \leq \rho_{j}$ for all $j$.

Form $n$ intervals

$$
D_{j}=\left[\widehat{r}_{j}(x)-2 \rho_{j}, \widehat{r}_{j}(x)+2 \rho_{j}\right], \quad j=1, \ldots, n
$$

The adaptive estimator is defined to be

$$
\begin{equation*}
\widehat{r}(x)=\widehat{r}_{j^{+}}(x) \tag{9.48}
\end{equation*}
$$

where $j_{+}$is the largest integer such that the intervals overlap:

$$
\begin{equation*}
j_{+}=\max \left\{k: \bigcap_{j=1}^{k} D_{j} \neq \emptyset\right\} . \tag{9.49}
\end{equation*}
$$

We now outline why $\widehat{r}(x)$ is adaptive.
Let $\bar{r}_{j}=\mathbb{E}\left(\widehat{r}_{j}\right)$ and note that

$$
\left|\widehat{r}_{j}-r(x)\right| \leq\left|\widehat{r}_{j}-\bar{r}_{j}\right|+\left|\bar{r}_{j}-r(x)\right|=b_{j}+\left|\xi_{j}\right|
$$

where $b_{j}=\left|\widehat{r}_{j}-\bar{r}_{j}\right|$ is the bias. Assume - as is the case for most smoothersthat $b_{j}$ decreases in $j$ while $s_{j}^{2}=\mathbb{V}\left(\xi_{j}\right)$ increases in $j$. Let

$$
j^{*}=\max \left\{j: b_{j} \leq \rho_{j}\right\}
$$

The bandwidth $h_{j^{*}}$ balances the bias and variance. The risk of the estimator with bandwidth $h_{j^{*}}$ is of order $\rho_{j^{*}}$. Thus, $\rho_{j^{*}}$ is the risk of an oracle estimator that knows the best bandwidth. We will refer to $\rho_{j^{*}}$ as the oracle risk.

For $j \leq j^{*},\left|\widehat{r}_{j}-r(x)\right| \leq b_{j}+\rho_{j} \leq 2 \rho_{j}$. Hence, $r(x) \in D_{j}$ for all $j \leq j^{*}$. In particular, all $D_{j}, j \leq j_{*}$ have at least one point in common, namely, $r(x)$. By the definition of $j^{+}$it follows that $j^{*} \leq j^{+}$. Also, every point in $D_{j^{*}}$ is at most $4 \rho_{j^{*}}$ away from $r(x)$. Finally, $D_{j^{+}} \bigcap D_{j^{*}} \neq \emptyset$ (from the definition of $j^{*}$ ) and $D_{j^{+}}$has half-length $2 \rho_{j^{+}}$. Thus,

$$
|\widehat{r}(x)-r(x)| \leq\left|\widehat{r}(x)-\widehat{r}_{j^{*}}\right|+\left|\widehat{r}_{j^{*}}-r(x)\right| \leq 2 \rho_{j^{+}}+4 \rho_{j^{*}} \leq 6 \rho_{j^{*}}
$$

We conclude that, with probability tending to $1,|\widehat{r}(x)-r(x)| \leq 6 \rho_{j^{*}}$.
A specific implementation of the idea is as follows. Fit a polynomial of order $m$ over an interval $\Delta$ containing the point $x \in(0,1)$. The resulting estimator is

$$
\widehat{r}_{\Delta}(x)=\sum_{i} \alpha_{\Delta}\left(x_{i}, x\right) Y_{i}
$$

for weights $\alpha_{\Delta}\left(x_{i}, x\right)$. The weights can be written as

$$
\alpha_{\Delta}(x, u)=\sum_{j=0}^{m} q_{\Delta}^{(j)}(x)\left(\frac{u-a}{b-a}\right)^{j}
$$

where $a=\min \left\{x_{i}: x_{i} \in \Delta\right\}$ and $b=\max \left\{x_{i}: x_{i} \in \Delta\right\}$. It can be shown that

$$
\left|\alpha_{\Delta}\left(x_{i}\right)\right| \leq \frac{c_{m}}{N_{\Delta}}
$$

where $N_{\delta}$ is the number of points in $\Delta$ and $c_{m}$ is a constant depending only on $m$. Also, it can be shown that the quantity

$$
\tau_{m} \equiv \frac{N_{\Delta}}{\theta_{m}} \max _{i, j}\left|q_{\Delta}^{(j)}\left(x_{i}\right)\right|
$$

depends only on $m$. Let

$$
D_{\Delta}=\left[\widehat{r}_{\Delta}-2 \kappa s_{\Delta}, \widehat{r}_{\Delta}+2 \kappa s_{\Delta}\right]
$$

where

$$
s_{\Delta}=\left(\sum_{i} \alpha_{\Delta}^{2}\left(x_{i}\right)\right)^{1 / 2}
$$

and

$$
\kappa_{n}=2 \sigma \sqrt{(m+2) \log n} 2^{m}(m+1) c_{m}^{2} \tau_{m}
$$

Now let $\widehat{\Delta}$ be the largest interval $\Delta$ containing $x$ such that

$$
\bigcap_{\substack{\delta \in \mathcal{D} \\ \delta \in \Delta}} D_{\delta} \neq \emptyset
$$

where $\mathcal{D}$ denotes all intervals containing $x$. Finally let $\widehat{r}(x)=\widehat{r}_{\widehat{\Delta}}(x)$.
Now we suppose that $r$ is $\ell$ times differentiable on an interval $\Delta_{0}=[x-$ $\left.\delta_{0}, x+\delta_{0}\right] \subset[0,1]$ for some $\ell \geq 1$. Also, assume that

$$
\left(\int_{\Delta_{0}}\left|r^{(\ell)}(t)\right|^{p} d t\right)^{1 / p} \leq L
$$

for some $L>0$, and $p \geq 1$, where

$$
\ell \leq m+1, \quad p \ell>1
$$

Then we have ${ }^{4}$ the following:
9.50 Theorem (Goldenshluger and Nemirovski 1997). Under the conditions above,

$$
\begin{equation*}
\left(\mathbb{E}|\widehat{r}(x)-r(x)|^{2}\right)^{1 / 2} \leq C B_{n} \tag{9.51}
\end{equation*}
$$

for some constant $C>0$ that only depends on $m$, where

$$
\begin{equation*}
B_{n}=\left[\left(\frac{\log n}{n}\right)^{(p \ell-1) /(2 p \ell+p-2)} L^{p /(2 p \ell+p-2)}+\sqrt{\frac{\log n}{n \delta_{0}}}\right] . \tag{9.52}
\end{equation*}
$$

The right-hand side of (9.51) is, except for logarithmic factors, the best possible risk. From the results of Lepskii (1991), the logarithmic factor is unavoidable. Since the estimator did not make use of the smoothness parameters $p, \ell$ and $L$, this means that the estimator adapts to the unknown smoothness.

Let us also briefly describe (a version of) the method of Lepski et al. (1997). Let $\mathcal{H}=\left\{h_{0}, h_{1}, \ldots, h_{m}\right\}$ be a set of bandwidths, where $h_{j}=a^{-j}$ where $a>1$

[^23](they use $a=1.02$ ) and $m$ is such that $h_{m} \approx \sigma^{2} / n$. Let $\widehat{r}_{h}(x)$ be the kernel estimator based on bandwidth $h$ and kernel $K$. For each bandwidth $h$, we test the hypothesis that further reducing $h$ does not significantly improve the fit. We take $\widehat{h}$ to be the largest bandwidth for which this test does not reject. Specifically,
$$
\widehat{h}=\max \left\{h \in \mathcal{H}:\left|\widehat{r}_{h}(x)-\widehat{r}_{\eta}(x)\right| \leq \psi(h, \eta), \quad \text { for all } \eta<h, \eta \in \mathcal{H}\right\}
$$
where
$$
\psi(h, \eta)=\frac{D \sigma}{n \eta} \sqrt{1+\log \left(\frac{1}{\eta}\right)}
$$
and $D>2(1+\|K\| \sqrt{14})$. They show that this bandwidth selection method yields to an estimator that adapts over a wide range of Besov spaces.

### 9.10 Do Adaptive Methods Work?

Let us take a closer look at the idea of spatial adaptation following some ideas in Donoho and Johnstone (1994). Let $A_{1}, \ldots, A_{L}$ be intervals that partition $[0,1]$ :

$$
A_{1}=\left[a_{0}, a_{1}\right), A_{2}=\left[a_{1}, a_{2}\right), \ldots, A_{L}=\left[a_{L-1}, a_{L}\right]
$$

where $a_{0}=0$ and $a_{L}=1$. Suppose that $r$ is a piecewise polynomial, so that $r(x)=\sum_{\ell=1}^{L} p_{\ell}(x) I\left(x \in A_{\ell}\right)$ where $p_{\ell}$ is a polynomial of degree $D$ on the set $A_{\ell}$. If the breakpoints $a=\left(a_{1}, \ldots, a_{L}\right)$ and the degree $D$ are known, then we can fit a polynomial of degree $D$ over each $A_{\ell}$ using least squares. This is a parametric problem and the risk is of order $O(1 / n)$. If we do not know the breakpoints $a$ and we fit a kernel regression then it can be shown that the risk is in general not better than order $O(1 / \sqrt{n})$ due to the possibility of discontinuities at the breakpoints. In contrast, Donoho and Johnstone show that the wavelet method has risk of order $O(\log n / n)$. This is an impressive theoretical achievement.

On the other hand, examples like those seen here suggest that the practical advantages are often quite modest. And from the point of view of inference (which features in the estimate are real?) the results in Chapter 7 show that wavelet confidence balls cannot converge faster than $n^{-1 / 4}$. To see this, consider the piecewise polynomial example again. Even knowing that $r$ is a piecewise polynomial, it still follows that the vector $\left(r\left(x_{1}\right), \ldots, r\left(x_{n}\right)\right)$ can take any value in $\mathbb{R}^{n}$. It then follows from Theorem 7.71 that no confidence ball can shrink faster than $n^{-1 / 4}$. Thus, we are in the peculiar situation that
the function estimate may converge quickly but the confidence set converges slowly.

So, do adaptive methods work or not? If one needs accurate function estimates and the noise level $\sigma$ is low, then the answer is that adaptive function estimators are very effective. But if we are facing a standard nonparametric regression problem and we are interested in confidence sets, then adaptive methods do not perform significantly better than other methods such as fixed bandwidth local regression.

### 9.11 Bibliographic Remarks

A good introduction to wavelets is Ogden (1997). A more advanced treatment can be found in Härdle et al. (1998). The theory of statistical estimation using wavelets has been developed by many authors, especially David Donoho and Iain Johnstone. The main ideas are in the following remarkable series of papers: Donoho and Johnstone (1994), Donoho and Johnstone (1995), Donoho et al. (1995), and Donoho and Johnstone (1998). The material on confidence sets is from Genovese and Wasserman (2005).

### 9.12 Appendix

Localization of Wavelets. The idea that wavelets are more localized than sines and cosines can be made precise. Given a function $f$, define its radius to be

$$
\Delta_{f}=\frac{1}{\|f\|}\left\{\int(x-\bar{x})^{2}|f(x)|^{2} d x\right\}^{1 / 2}
$$

where

$$
\bar{x}=\frac{1}{\|f\|^{2}} \int x|f(x)|^{2} d x .
$$

Imagine drawing a rectangle on the plane with sides of length $\Delta_{f}$ and $\Delta_{f^{*}}$. For a function like a cosine, this is a rectangle with 0 width on the $y$-axis and infinite width in the $x$-axis. We say that a cosine is localized in frequency but nonlocalized in space. In contrast, a wavelet has a rectangle with finite length in both dimensions. Hence, wavelets are localized in frequency and space. There is a limit to how well we can simultaneously localize in frequency and space.
9.53 Theorem (Heisenberg's uncertainty relation). We have that

$$
\begin{equation*}
\Delta_{f} \Delta_{f^{*}} \geq \frac{1}{2} \tag{9.54}
\end{equation*}
$$

with equality when $f$ is a Normal density.
This inequality is called the Heisenberg uncertainty principle since it first appeared in the physics literature when Heisenberg was developing quantum mechanics.

Fast Computations for Wavelets. The scaling coefficients make it easy to compute the wavelet coefficients. Recall that

$$
\alpha_{j k}=\int f(x) \phi_{j k}(x) d x \text { and } \beta_{j k}=\int f(x) \psi_{j k}(x) d x
$$

By definition, $\phi_{j k}(x)=2^{j / 2} \phi\left(2^{j} x-k\right)$ and by (9.16), $\phi\left(2^{j} x-k\right)=\sum_{r} \ell_{r} \phi_{1, r}\left(2^{j} x-\right.$ k). Hence,

$$
\begin{aligned}
\phi_{j k}(x) & =\sum_{r} \ell_{r} 2^{j / 2} \phi_{1, r}\left(2^{j} x-k\right)=\sum_{r} \ell_{r} 2^{(j+1) / 2} \phi\left(2^{j+1} x-2 k-r\right) \\
& =\sum_{r} \ell_{r} \phi_{j+1, \ell+2 k}(x)=\sum_{r} \ell_{r-2 k} \phi_{j+1, r}(x) .
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\alpha_{j k} & =\int f(x) \phi_{j k}(x) d x=\int f(x) \sum_{r} \ell_{r-2 k} \phi_{j+1, r}(x) d x \\
& =\sum_{r} \ell_{r-2 k} \int f(x) \phi_{j+1, r}(x) d x=\sum_{r} \ell_{r-2 k} \alpha_{j+1, r} .
\end{aligned}
$$

By similar calculations, one can show that $\beta_{j k}=\sum_{r}(-1)^{r+1} \ell_{-r+2 k+1} \alpha_{j+1, r}$. In summary we get the following cascade equations:

|  | The Cascade Equations |
| ---: | :--- |
| $\alpha_{j k}$ | $=\sum_{r} \ell_{r-2 k} \alpha_{j+1, r}$ |
| $\beta_{j k}$ | $=\sum_{r}(-1)^{r+1} \ell_{-r+2 k+1} \alpha_{j+1, r}$. |

Once we have the scaling coefficients $\left\{\alpha_{j k}\right\}$ for some $J$, we may determine $\left\{\alpha_{j k}\right\}$ and $\left\{\beta_{j k}\right\}$ for all $j<J$ by way of the cascade equations (9.55) and


FIGURE 9.10. The cascade algorithm.
(9.56). This method of computing the coefficients is called the pyramid algorithm or the cascade algorithm.

In regression problems, we will use the data $Y_{1}, \ldots, Y_{n}$ to approximate the scaling coefficients at some high-level $J$. The other coefficients are found from the pyramid algorithm. This process is called the discrete wavelet transform. It only requires $O(n)$ operations for $n$ data points.

These ideas can be expressed in the language of signals and filters. A signal is defined to be a sequence $\left\{f_{k}\right\}_{k \in \mathbb{Z}}$ such that $\sum_{k} f_{k}^{2}<\infty$. A filter $A$ is a function that operates on signals. A filter $A$ can be represented by some coefficients $\left\{a_{k}\right\}_{k \in \mathbb{Z}}$ and the action of $A$ on a signal $f$-called discrete convolution-produces a new signal denoted by $A f$ whose $k^{\text {th }}$ coefficient is

$$
\begin{equation*}
(A f)_{k}=\sum_{r} a_{r-2 k} f_{r} \tag{9.57}
\end{equation*}
$$

Let $\alpha_{j}=\left\{\alpha_{j k}\right\}_{k \in \mathbb{Z}}$ be the scaling coefficients at level $j$. Let $L$ be the filter with coefficients $\left\{\ell_{k}\right\}$. $L$ is called a low-pass filter. From Equation (9.55) we get

$$
\begin{equation*}
\alpha_{j-1}=L \alpha_{j,} \quad \text { and } \quad \alpha_{j-m}=L^{m} \alpha_{j} \tag{9.58}
\end{equation*}
$$

where $L^{m}$ means: apply the filter $m$ times. Let $H$ be the filter with coefficients $h_{k}=(-1)^{k+1} \ell_{1-k}$. Then (9.56) implies that

$$
\begin{equation*}
\beta_{j-1}=H \alpha_{j,} \quad \text { and } \quad \beta_{j-m},=H L^{m-1} \alpha_{j,} . \tag{9.59}
\end{equation*}
$$

$H$ is called a high-pass filter. See Figure 9.10 for a schematic representation of the algorithm.

### 9.13 Exercises

1. Show that, for the Haar wavelet, $\ell_{0}=\ell_{1}=2^{-1 / 2}$ and $\ell_{k}=0$ for $k \neq 0,1$.
2. Prove Theorem 9.6.
3. Prove that the Haar wavelets form an MRA; see Example 9.15.
4. Prove equation (9.17).
5. Generate data

$$
Y_{i}=r(i / n)+\sigma \epsilon_{i}
$$

where $r$ is the doppler function, $n=1024$, and $\epsilon_{i} \sim N(0,1)$.
(a) Use wavelets to fit the curve. Try the following shrinkage methods:
(i) James-Stein applied to each resolution level, (ii) universal shrinkage and (iii) SureShrink. Try $\sigma=.01, \sigma=.1, \sigma=1$. Compare the function estimates to the REACT method and to local linear regression.
(b) Repeat (a) but add outliers to the $\epsilon_{i}$ by generating them from:

$$
\epsilon_{i} \sim .95 N(0,1)+.05 N(0,4)
$$

How does this affect your results?
6 . Let $X_{1}, \ldots, X_{n} \sim f$ for some density $f$ on $[0,1]$. Let's consider constructing a wavelet histogram. Let $\phi$ and $\psi$ be the Haar father and mother wavelet. Write

$$
f(x) \approx \phi(x)+\sum_{j=0}^{J} \sum_{k=0}^{2^{j}-1} \beta_{j k} \psi_{j k}(x)
$$

where $J \approx \log _{2}(n+1)$. The total number of wavelet coefficients is about $n$. Now,

$$
\beta_{j k}=\int_{0}^{1} \psi_{j k}(x) f(x) d x=E_{f}\left[\psi_{j k}(X)\right]
$$

An unbiased estimator of $\beta_{j k}$ is

$$
\widetilde{\beta}_{j k}=\frac{1}{n} \sum_{i=1}^{n} \psi_{j k}\left(X_{i}\right)
$$

(a) For $x<y$ define

$$
N_{x, y}=\sum_{i=1}^{n} I\left(x \leq X_{i}<y\right)
$$

Show that

$$
\widetilde{\beta}_{j k}=\frac{2^{j / 2}}{n}\left(N_{2 k /\left(2^{j+1}\right),(2 k+1) /\left(2^{j+1}\right)}-N_{(2 k+1) /\left(2^{j+1}\right),(2 k+2) /\left(2^{j+1}\right)}\right) .
$$

(b) Show that

$$
\widetilde{\beta}_{j k} \approx N\left(\beta_{j k}, \sigma_{j k}^{2}\right)
$$

Find an expression for $\sigma_{j k}$.
(c) Consider the shrinkage estimator $\widehat{\beta}_{j k}=a_{j} \widetilde{\beta}_{j k}$. (Same shrinkage coefficient across each level of wavelets.) Take $\sigma_{j k}$ as known, and find $a_{j}$ to minimize Stein's unbiased risk estimate.
(d) Now find an estimate of $\sigma_{j k}$. Insert this estimate into your formula for $a_{j}$. You now have a method for estimating the density. Your final estimate is

$$
\widehat{f}(x)=\phi(x)+\sum_{j} \sum_{k} \widehat{\beta}_{j k} \psi_{j k}(x) .
$$

Try this on the geyser duration data from the book website. Compare it to the histogram with binwidth chosen using cross-validation.
(e) Notice that $\widetilde{\beta}_{j k}$ is $2^{j / 2}$ times the difference of two sample proportions. Furthermore, $\beta_{j k}=0$ corresponds to the two underlying Binomial parameters being equal. Hence, we can form a hard threshold estimator by testing (at some level $\alpha$ ) whether each $\beta_{j k}=0$ and only keeping those that are rejected. Try this for the data in part (d). Try different values of $\alpha$.
7. Let $R_{n}^{*}$ be the oracle risk defined in equation (9.35).
(a) Show that

$$
R_{n}^{*}=\sum_{i=1}^{n}\left(\frac{\sigma^{2}}{n} \wedge \theta_{i}^{2}\right)
$$

where $a \wedge b=\min \{a, b\}$.
(b) Compare $R_{n}^{*}$ to the Pinsker bound in Theorem 7.28. For which vectors $\theta \in \mathbb{R}^{n}$ is $R_{n}^{*}$ smaller than than the Pinsker bound?
8. Find an exact expression for the risk of the hard threshold rule and the soft threshold rule (treating $\sigma$ as known).
9. Generate data

$$
Y_{i}=r(i / n)+\sigma \epsilon_{i}
$$

where $r$ is the Doppler function, $n=1024$, and $\epsilon_{i} \sim N(0,1)$ and $\sigma=0.1$. Apply the ICI method from Section 9.9 to estimate $r$. Use a kernel estimator and take the grid of bandwidths to be $\{1 / n, \ldots, 1\}$. First, take $\sigma$ as known. Then estimate $\sigma$ using one of the methods from Section 5.6. Now apply the method of Lepski et al. (1997) also from Section 5.6.

## 10

## Other Topics

In this chapter we mention some other issues related to nonparametric inference including: measurement error, inverse problems, nonparametric Bayesian inference, semiparametric inference, correlated errors, classification, sieves, shape-restricted inference, testing, and computation.

### 10.1 Measurement Error

Suppose we are interested in regressing the outcome $Y$ on a covariate $X$ but we cannot observe $X$ directly. Rather, we observe $X$ plus noise $U$. The observed data are $\left(X_{1}^{\bullet}, Y_{1}\right), \ldots,\left(X_{n}^{\bullet}, Y_{n}\right)$ where

$$
\begin{aligned}
Y_{i} & =r\left(X_{i}\right)+\epsilon_{i} \\
X_{i}^{\bullet} & =X_{i}+U_{i}, \quad \mathbb{E}\left(U_{i}\right)=0 .
\end{aligned}
$$

This is called a measurement error problem or an errors-in-variables problem. A good reference is Carroll et al. (1995) which we follow closely in this section.

The model is illustrated by the directed graph in Figure 10.1. It is tempting to ignore the error and just regress $Y$ on $X^{\bullet}$ but this leads to inconsistent estimates of $r(x)$.


FIGURE 10.1. Regression with measurement error. $X$ is circled to show that it is not observed. $X^{\bullet}$ is a noisy version of $X$. If you regress $Y$ on $X^{\bullet}$, you will get an inconsistent estimate of $r(x)$.

Before discussing the nonparametric problem, we first consider the linear regression version of this problem. The model is

$$
\begin{aligned}
Y_{i} & =\beta_{0}+\beta_{1} X_{i}+\epsilon_{i} \\
X_{i}^{\bullet} & =X_{i}+U_{i}
\end{aligned}
$$

Let $\sigma_{x}^{2}=\mathbb{V}(X)$, and assume that $\epsilon$ is independent of $X$, has mean 0 and variance $\sigma_{\epsilon}^{2}$. Also assume that $U$ is independent of $X$, with mean 0 and variance $\sigma_{u}^{2}$. Let $\widehat{\beta}_{1}$ be the least squares estimator of $\beta_{1}$ obtained by regressing the $Y_{i}$ 's on the $X_{i}^{\bullet}$ 's. It can be shown (see Exercise 2) that

$$
\begin{equation*}
\widehat{\beta} \xrightarrow{\text { a.s. }} \lambda \beta_{1} \tag{10.1}
\end{equation*}
$$

where

$$
\begin{equation*}
\lambda=\frac{\sigma_{x}^{2}}{\sigma_{x}^{2}+\sigma_{u}^{2}}<1 \tag{10.2}
\end{equation*}
$$

Thus, the effect of the measurement is to bias the estimated slope towards 0 , an effect that is usually called attenuation bias. Staudenmayer and Ruppert (2004) showed that a similar result holds for nonparametric regression. If we use local polynomial regression (with odd degree polynomial) and do not account for the measurement error, the estimator $\widehat{r}_{n}(x)$ has, asymptotically, an excess bias of

$$
\begin{equation*}
\sigma_{u}^{2}\left(\frac{f^{\prime}(x)}{f(x)} r^{\prime}(x)+\frac{r^{\prime \prime}(x)}{2}\right) \tag{10.3}
\end{equation*}
$$

where $f$ is the density of $X$.
Returning to the linear case, if there are several observed values of $X^{\bullet}$ for each $X$ then $\sigma_{u}^{2}$ can be estimated. Otherwise, $\sigma_{u}^{2}$ must be estimated by external means such as through background knowledge of the noise mechanism. For
our purposes, we will assume that $\sigma_{u}^{2}$ is known. Since, $\sigma_{\bullet}^{2}=\sigma_{x}^{2}+\sigma_{u}^{2}$, we can estimate $\sigma_{x}^{2}$ by

$$
\begin{equation*}
\widehat{\sigma}_{x}^{2}=\widehat{\sigma}_{\bullet}^{2}-\sigma_{u}^{2} \tag{10.4}
\end{equation*}
$$

where $\widehat{\sigma}_{\bullet}^{2}$ is the sample variance of the $X_{i}^{\bullet}$. Plugging these estimates into (10.2), we get an estimate $\widehat{\lambda}=\left(\widehat{\sigma}_{\bullet}^{2}-\sigma_{u}^{2}\right) / \widehat{\sigma}_{\bullet}^{2}$ of $\lambda$. An estimate of $\beta_{1}$ is

$$
\begin{equation*}
\widetilde{\beta}_{1}=\frac{\widehat{\beta}_{1}}{\widehat{\lambda}}=\frac{\widehat{\sigma}_{\bullet}^{2}}{\widehat{\sigma}_{\bullet}^{2}-\sigma_{u}^{2}} \widehat{\beta}_{1} \tag{10.5}
\end{equation*}
$$

This is called the method of moments estimator. See Fuller (1987) for more details.

Another method for correcting the attenuation bias is SIMEX which stands for simulation extrapolation and is due to Cook and Stefanski (1994) and Stefanski and Cook (1995). Recall that the least squares estimate $\widehat{\beta}_{1}$ is a consistent estimate of

$$
\frac{\beta_{1} \sigma_{x}^{2}}{\sigma_{x}^{2}+\sigma_{u}^{2}}
$$

Generate new random variables

$$
\widetilde{X}_{i}=X_{i}^{\bullet}+\sqrt{\rho} U_{i}
$$

where $U_{i} \sim N(0,1)$. The least squares estimate obtained by regressing the $Y_{i}$ 's on the $\widetilde{X}_{i}$ 's is a consistent estimate of

$$
\begin{equation*}
\Omega(\rho)=\frac{\beta_{1} \sigma_{x}^{2}}{\sigma_{x}^{2}+(1+\rho) \sigma_{u}^{2}} \tag{10.6}
\end{equation*}
$$

Repeat this process $B$ times (where $B$ is large) and denote the resulting estimators by $\widehat{\beta}_{1,1}(\rho), \ldots, \widehat{\beta}_{1, B}(\rho)$. Then define

$$
\widehat{\Omega}(\rho)=\frac{1}{B} \sum_{b=1}^{B} \widehat{\beta}_{1, b}(\rho)
$$

Now comes some clever sleight of hand. Setting $\rho=-1$ in (10.6) we see that $\Omega(-1)=\beta_{1}$ which is the quantity we want to estimate. The idea is to compute $\widehat{\Omega}(\rho)$ for a range of values of $\rho$ such as $0,0.5,1.0,1.5,2.0$. We then extrapolate the curve $\widehat{\Omega}(\rho)$ back to $\rho=-1$; see Figure 10.2. To do the extrapolation, we fit the values $\widehat{\Omega}\left(\rho_{j}\right)$ to the curve

$$
\begin{equation*}
G\left(\rho ; \gamma_{1}, \gamma_{2}, \gamma_{3}\right)=\gamma_{1}+\frac{\gamma_{2}}{\gamma_{3}+\rho} \tag{10.7}
\end{equation*}
$$

using standard nonlinear regression. Once we have estimates of the $\gamma$ 's, we take

$$
\begin{equation*}
\widetilde{\beta}_{1}=G\left(-1 ; \widehat{\gamma}_{1}, \widehat{\gamma}_{2}, \widehat{\gamma}_{3}\right) \tag{10.8}
\end{equation*}
$$



FIGURE 10.2. In the SIMEX method we extrapolate $\widehat{\Omega}(\rho)$ back to $\rho=-1$.
as our corrected estimate of $\beta_{1}$. Fitting the nonlinear regression (10.7) is inconvenient; it often suffices to approximate $G(\rho)$ with a quadratic. Thus, we fit the $\widehat{\Omega}\left(\rho_{j}\right)$ values to the curve

$$
Q\left(\rho ; \gamma_{1}, \gamma_{2}, \gamma_{3}\right)=\gamma_{1}+\gamma_{2} \rho+\gamma_{3} \rho^{2}
$$

and the corrected estimate of $\beta_{1}$ is

$$
\widetilde{\beta}_{1}=Q\left(-1 ; \widehat{\gamma}_{1}, \widehat{\gamma}_{2}, \widehat{\gamma}_{3}\right)=\widehat{\gamma}_{1}-\widehat{\gamma}_{2}+\widehat{\gamma}_{3} .
$$

An advantage of SIMEX is that it extends readily to nonparametric regression. Let $\widehat{r}_{n}(x)$ be an uncorrected estimate of $r(x)$ obtained by regressing the $Y_{i}$ 's on the $X_{i}{ }^{\mathbf{0}}$ 's in the nonparametric problem

$$
\begin{aligned}
Y_{i} & =r\left(X_{i}\right)+\epsilon_{i} \\
X_{i}^{\bullet} & =X_{i}+U_{i} .
\end{aligned}
$$

Now perform the SIMEX algorithm to get $\widehat{r}_{n}(x, \rho)$ and define the corrected estimator $\widetilde{r}_{n}(x)=\widehat{r}_{n}(x,-1)$. There remains the problem of choosing the smoothing parameter. This is an area of active research; see Staudenmayer and Ruppert (2004), for example.

A more direct way to deal with measurement error is suggested by Fan and Truong (1993). They propose the kernel estimator

$$
\begin{equation*}
\widehat{r}_{n}(x)=\frac{\sum_{i=1}^{n} K_{n}\left(\frac{x-X_{i}^{\bullet}}{h_{n}}\right) Y_{i}}{\sum_{i=1}^{n} K_{n}\left(\frac{x-X_{i}^{\bullet}}{h_{n}}\right)} \tag{10.9}
\end{equation*}
$$

where

$$
K_{n}(x)=\frac{1}{2 \pi} \int e^{-i t x} \frac{\phi_{K}(t)}{\phi_{U}\left(t / h_{n}\right)} d t
$$

where $\phi_{K}$ is the Fourier transform of a kernel $K$ and $\phi_{U}$ is the characteristic function of $U$. This is a standard kernel estimator except for the unusual kernel $K_{n}$, which will be motivated later in this section, after equation (10.21).

Yet another way to deal with measurement error, due to Stefanski (1985), is motivated by considering the asymptotics in which $\sigma_{u} \rightarrow 0$ rather than keeping $\sigma_{u}$ fixed as $n$ increases. Let us apply Stefanski's "small $\sigma_{u}$ " approach to the nonparametric regression problem. Write the uncorrected estimator as

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}\left(x, X_{i}^{\bullet}\right) \tag{10.10}
\end{equation*}
$$

where we have written the weights as $\ell_{i}\left(x, X_{i}^{\bullet}\right)$ to emphasize the dependence on $X_{i}^{\bullet}$. If the $X_{i}$ 's had been observed, the estimator of $r$ would be

$$
r_{n}^{*}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}\left(x, X_{i}\right)
$$

Expanding $\ell_{i}\left(x, X_{i}^{\bullet}\right)$ around $X_{i}$ we have

$$
\begin{equation*}
\widehat{r}_{n}(x) \approx r_{n}^{*}(x)+\sum_{i=1}^{n} Y_{i}\left(X_{i}^{\bullet}-X_{i}\right) \ell^{\prime}\left(x, X_{i}\right)+\frac{1}{2} \sum_{i=1}^{n} Y_{i}\left(X_{i}^{\bullet}-X_{i}\right)^{2} \ell^{\prime \prime}\left(x, X_{i}\right) \tag{10.11}
\end{equation*}
$$

Taking expectations, we see that the excess bias due to measurement error (conditional on the $X_{i} \mathrm{~s}$ ) is

$$
\begin{equation*}
b(x)=\frac{\sigma_{u}^{2}}{2} \sum_{i=1}^{n} r\left(X_{i}\right) \ell^{\prime \prime}\left(x, X_{i}\right) \tag{10.12}
\end{equation*}
$$

We can estimate $b(x)$ with

$$
\begin{equation*}
\widehat{b}(x)=\frac{\sigma_{u}^{2}}{2} \sum_{i=1}^{n} \widehat{r}\left(X_{i}^{\bullet}\right) \ell^{\prime \prime}\left(x, X_{i}^{\bullet}\right) \tag{10.13}
\end{equation*}
$$

This yields a bias corrected estimate for $r$, namely,

$$
\begin{equation*}
\widetilde{r}_{n}(x)=\widehat{r}_{n}(x)-\widehat{b}(x)=\widehat{r}_{n}(x)-\frac{\sigma_{u}^{2}}{2} \sum_{i=1}^{n} \widehat{r}\left(X_{i}^{\bullet}\right) \ell^{\prime \prime}\left(x, X_{i}^{\bullet}\right) . \tag{10.14}
\end{equation*}
$$

This estimator still has measurement error bias due to the estimation of $b(x)$ but, for small $\sigma_{u}^{2}$, it will be less biased than $\widehat{r}_{n}$.

Now consider density estimation. Suppose that

$$
\begin{aligned}
X_{1}, \ldots, X_{n} & \sim F \\
X_{i}^{\bullet} & =X_{i}+U_{i}, \quad i=1, \ldots, n
\end{aligned}
$$

where, as before, the $X_{i}$ 's are not observed. We want to estimate the density $f(x)=F^{\prime}(x)$. The density of $X^{\bullet}$ is

$$
\begin{equation*}
f^{\bullet}(x)=\int f(s) f_{U}(x-s) d s \tag{10.15}
\end{equation*}
$$

where $f_{U}$ is the density of $U$. Since $f^{\bullet}$ is the convolution of $f$ and $f_{U}$, the problem of estimating $f$ is called deconvolution.

One way to estimate $f$ is to use Fourier inversion. Let $\psi(t)=\int e^{i t x} f(x) d x$ denote the Fourier transform (characteristic function) of $X$ and define $\psi^{\bullet}$ and $\psi_{U}$ similarly. Since $X^{\bullet}=X+U$, it follows that $\psi^{\bullet}(t)=\psi(t) \psi_{U}(t)$ and hence

$$
\begin{equation*}
\psi(t)=\frac{\psi^{\bullet}(t)}{\psi_{U}(t)} \tag{10.16}
\end{equation*}
$$

If $\widehat{f}^{\bullet}$ is an estimate of $f^{\bullet}$ then

$$
\begin{equation*}
\widehat{\psi}^{\bullet}(t)=\int e^{i t x} \widehat{f}^{\bullet}(x) d x \tag{10.17}
\end{equation*}
$$

is an estimate of $\psi^{\bullet}$. By Fourier inversion and equation (10.16),

$$
\begin{equation*}
f(x)=\frac{1}{2 \pi} \int e^{-i t x} \psi(t) d t=\frac{1}{2 \pi} \int e^{-i t x} \frac{\psi^{\bullet}(t)}{\psi_{U}(t)} d t \tag{10.18}
\end{equation*}
$$

which suggests the estimator

$$
\begin{equation*}
\widehat{f}(x)=\frac{1}{2 \pi} \int e^{-i t x} \frac{\widehat{\psi}^{\bullet}(t)}{\psi_{U}(t)} d t \tag{10.19}
\end{equation*}
$$

In particular, if $\widehat{f}^{\bullet}$ is a kernel estimator,

$$
\widehat{f}^{\bullet}(x)=\frac{1}{n h} \sum_{j=1}^{n} K\left(\frac{x-X_{j}^{\bullet}}{h}\right)
$$

then $\widehat{f}(x)$ can be written as

$$
\begin{equation*}
\widehat{f}(x)=\frac{1}{n h} \sum_{j=1}^{n} K_{*}\left(\frac{x-X_{j}^{\bullet}}{h}, h\right) \tag{10.20}
\end{equation*}
$$

where

$$
\begin{equation*}
K_{*}(t, h)=\frac{1}{2 \pi} \int e^{i t u} \frac{\psi_{K}(u)}{\psi_{U}(u / h)} d u \tag{10.21}
\end{equation*}
$$

and $\psi_{K}$ is the Fourier transform of the kernel $K$. Equations (10.20) and (10.21) motivate the kernel regression estimator (10.9).

The risk of (10.20) is

$$
\begin{equation*}
\mathbb{E}\left(\widehat{f}_{Y}(y)-f(y)\right)^{2} \approx c h^{4}+\frac{1}{2 \pi n h} \int\left(\frac{\psi_{K}(t)}{\left|\psi_{U}(t / h)\right|}\right)^{2} d t \tag{10.22}
\end{equation*}
$$

where

$$
\begin{equation*}
c=\frac{1}{4} \int y^{2} K(y) d y \int\left(f^{\prime \prime}(y)\right)^{2} d y \tag{10.23}
\end{equation*}
$$

Note that $\psi_{U}(t / h)$ appears in the denominator of (10.22). Thus, if $\psi_{U}(t / h)$ has thin tails, the risk will be large. Now $\psi_{U}(t / h)$ has thin tails when $f_{U}$ is smooth, implying that if $f_{U}$ is smooth, the rate of convergence is slow. In particular, if $f_{U}$ is Normal, it can be shown that the best rate of convergence is $O(1 /(\log n))^{2}$ which is very slow. Stefanski (1990) showed the surprising result that, under fairly general conditions, the optimal bandwidth is $h=\sigma_{u} / \sqrt{\log n}$ which does not depend on $f$.

In these asymptotic calculations, $n$ is increasing while $\sigma_{u}^{2}=\mathbb{V}(U)$ stays fixed. As we mentioned earlier, a more realistic asymptotic calculation might also have $\sigma_{u}^{2}$ tending to 0 . In this approach, the convergence rate is less discouraging. The small $\sigma_{u}$ approach suggests the corrected estimator

$$
\widehat{f}_{n}(x)-\frac{\sigma_{u}^{2}}{2 n h} \sum_{i=1}^{n} K^{\prime \prime}\left(\frac{X_{\dot{j}}^{\bullet}-x}{h}\right)
$$

where $\widehat{f}_{n}$ is the naive, uncorrected kernel estimator using the $X_{i}^{\bullet} \mathrm{s}$.

### 10.2 Inverse Problems

A class of problems that are very similar to measurement error are inverse problems. In general, an inverse problem refers to reconstructing properties of an object given only partial information on an object. An example is trying to estimate the three-dimensional structure of an object given information from two-dimensional slices of that object. This is common in certain types of medical diagnostic procedures. Another example is image reconstruction when the image is subject to blurring. Our summary here follows O'Sullivan (1986).

In the regression framework, a common inverse problem is of the form

$$
\begin{equation*}
Y_{i}=T_{i}(r)+\epsilon_{i}, \quad i=1, \ldots, n \tag{10.24}
\end{equation*}
$$

where $r$ is the regression function of interest and $T_{i}$ is some operator acting on $r$. A concrete example, which we shall use for the rest of this section, is $T_{i}(r)=$
$\int K_{i}(s) r(s) d s$ for some smooth function $K_{i}$ such as $K_{i}(s)=e^{-\left(s-x_{i}\right)^{2} / 2}$. The model becomes

$$
\begin{equation*}
Y_{i}=\int K_{i}(s) r(s) d s+\epsilon_{i} \tag{10.25}
\end{equation*}
$$

If $K_{i}$ is a delta function at $x_{i}$, then (10.25) becomes the usual nonparametric regression model $Y_{i}=r\left(x_{i}\right)+\epsilon_{i}$. Think of $\int K_{i}(s) r(s) d s$ as a blurred version of $r$. There are now two types of information loss: the noise $\epsilon_{i}$ and blurring.

Suppose we estimate $r$ with a linear smoother as defined in Section 5.2:

$$
\begin{equation*}
\widehat{r}_{n}(x)=\sum_{i=1}^{n} Y_{i} \ell_{i}(x) \tag{10.26}
\end{equation*}
$$

The variance is the same as without blurring, namely, $\mathbb{V}\left(\widehat{r}_{n}(x)\right)=\sigma^{2} \sum_{i=1}^{n} \ell_{i}^{2}(x)$ but the mean has a different form:

$$
\mathbb{E}\left(\widehat{r}_{n}(x)\right)=\sum_{i=1}^{n} \ell_{i}(x) \int K_{i}(s) r(s) d s=\int A(x, s) r(s) d s
$$

where

$$
\begin{equation*}
A(x, s)=\sum_{i=1}^{n} \ell_{i}(x) K_{i}(s) \tag{10.27}
\end{equation*}
$$

is called the Backus-Gilbert averaging kernel.
Suppose that $r$ can be approximated as an expansion in some basis $\left\{\phi_{1}, \ldots, \phi_{k}\right\}$ (see Chapter 8), that is, $r(x)=\sum_{j=1}^{k} \theta_{j} \phi_{j}(x)$. Then,

$$
\int K_{i}(s) r(s) d s=\int K_{i}(s) \sum_{j=1}^{k} \theta_{j} \phi_{j}(s) d s=Z_{i}^{T} \theta
$$

where $\theta=\left(\theta_{1}, \ldots, \theta_{k}\right)^{T}$ and

$$
Z_{i}=\left(\begin{array}{c}
\int K_{i}(s) \phi_{1}(s) d s \\
\int K_{i}(s) \phi_{2}(s) d s \\
\vdots \\
\int K_{i}(s) \phi_{k}(s) d s
\end{array}\right)
$$

The model (10.25) can then be written as

$$
\begin{equation*}
Y=Z \theta+\epsilon \tag{10.28}
\end{equation*}
$$

where $Z$ is an $n \times k$ matrix with $i^{\text {th }}$ row equal to $Z_{i}^{T}, Y=\left(Y_{1}, \ldots, Y_{n}\right)^{T}$ and $\epsilon=\left(\epsilon_{1}, \ldots, \epsilon_{n}\right)^{T}$. It is tempting to estimate $\theta$ by the least squares estimator $\left(Z^{T} Z\right)^{-1} Z^{T} Y$. This may fail since $Z^{T} Z$ is typically not invertible in which
case the problem is said to be ill-posed. Indeed, this is a hallmark of inverse problems and corresponds to the fact that the function $r$ cannot be recovered, even in the absence of noise, due to the information loss incurred by blurring. Instead, it is common to use a regularized estimator such as $\widehat{\theta}=L Y$ where

$$
L=\left(Z^{T} Z+\lambda I\right)^{-1} Z^{T}
$$

$I$ is the identity matrix and $\lambda>0$ is a smoothing parameter that can be chosen by cross-validation. It should be noted that cross-validation is estimating the prediction error $n^{-1} \sum_{i=1}^{n}\left(\int K_{i}(s) r(s) d s-\int K_{i}(s) \widehat{r}(s) d s\right)^{2}$ rather than $\int(r(s)-\widehat{r}(s))^{2} d s$. In Chapter 5 we observed that these two loss functions were essentially the same. This is no longer true in the present context. In principle, it is also possible to devise a cross-validation estimator for the loss $\int(r(s)-\widehat{r}(s))^{2} d s$ but this estimator may be very unstable.

### 10.3 Nonparametric Bayes

Throughout this book we have taken a frequentist approach to inference. It is also possible to take a Bayesian approach. ${ }^{1}$ Indeed, Bayesian nonparametric inference is a thriving research enterprise in statistics as well as in machine learning. Good references are Ghosh and Ramamoorthi (2003), Dey et al. (1998), Walker (2004) and references therein. However, the area is far to large and growing far too quickly for us to cover the topic here.

A small sampling of relevant references, in addition to those already mentioned, includes Schwartz (1965), Diaconis and Freedman (1986), Barron et al. (1999b), Ghosal et al. (1999), Walker and Hjort (2001), Hjort (2003), Ghosal et al. (2000), Shen and Wasserman (2001), Zhao (2000), Huang (2004), Cox (1993), Freedman (1999), McAuliffe et al. (2004), Teh et al. (2004), Blei et al. (2004), Blei and Jordan (2004), and Wasserman (1998).

### 10.4 Semiparametric Inference

As the name suggests, semiparametric models are models that are part parametric and part nonparametric. An example is a partially linear regression model of the form

$$
\begin{equation*}
Y=\beta X+r(Z)+\epsilon \tag{10.29}
\end{equation*}
$$

[^24]where $r$ is some smooth function. The theory of inference for such models can get quite complex. Consider estimating $\beta$ in (10.29), for example. One strategy is to regress $Y_{i}-\widehat{r}_{n}\left(Z_{i}\right)$ on $X_{i}$ where $\widehat{r}_{n}$ is an estimate of $r$. Under appropriate conditions, if $\widehat{r}_{n}$ is chosen carefully, this will lead to good estimates of $\beta$. See Bickel et al. (1993) and Chapter 25 of van der Vaart (1998) for details.

### 10.5 Correlated Errors

If the errors $\epsilon_{i}$ in the regression model $Y_{i}=r\left(x_{i}\right)+\epsilon_{i}$ are correlated, then the usual methods can break down. In particular, positive correlation can make methods like cross-validation to choose very small bandwidths. There are several approaches for dealing with the correlation. In modified crossvalidation, we drop out blocks of observations instead of single observations. In partitioned cross-validation, we partition the data, and use one observation per partition to construct the cross-validation estimate. The estimates obtained this way are then averaged. The properties of these methods are discussed in Chu and Marron (1991). A review of nonparametric regression with correlated observations can be found in Opsomer et al. (2001).

### 10.6 Classification

In the classification problem we have data $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$ where $Y_{i}$ is discrete. We want to find a function $\widehat{h}$ so that, given a new $X$, we can predict $Y$ by $\widehat{h}(X)$. This is just like regression except for two things: (i) the outcome is discrete and (ii) we do not need to estimate the relationship between $Y$ and $X$ well; we only need to predict well.

In an earlier version of this book, there was a long chapter on classification. The topic is so vast that the chapter took on a life of its own and I decided to eliminate it. There are many good books on classification such as Hastie et al. (2001). I will only make a few brief comments here.

Suppose that $Y_{i} \in\{0,1\}$ is binary. A classifier is a function $h$ that maps each $x$ into $\{0,1\}$. A commonly used risk function for classification is $L(h)=$ $\mathbb{P}(Y \neq h(X))$. It can be shown that the optimal classification rule - called the

Bayes rule ${ }^{2}$-is

$$
h(x)= \begin{cases}1 & \text { if } r(x) \geq 1 / 2 \\ 0 & \text { if } r(x)<1 / 2\end{cases}
$$

where $r(x)=\mathbb{E}(Y \mid X=x)$. This suggests a natural (and not uncommon) approach to classification. Form an estimate $\widehat{r}_{n}(x)$ of $r(x)$ then estimate $h$ by

$$
\widehat{h}(x)= \begin{cases}1 & \text { if } \widehat{r}_{n}(x) \geq 1 / 2 \\ 0 & \text { if } \widehat{r}_{n}(x)<1 / 2 .\end{cases}
$$

Now even if $\widehat{r}_{n}$ is a poor estimator of $r, \widehat{h}$ might still be a good classifier. For example, if $r(x)=.6$ but $\widehat{r}_{n}(x)=.9$, we still have that $h(x)=\widehat{h}(x)=1$.

### 10.7 Sieves

A sieve is a sequence of models, indexed by sample size $n$, that increase in complexity as $n \rightarrow \infty$. A simple example is polynomial regression where the maximum degree of the polynomial $p(n)$ increases with $n$. Choosing $p(n)$ is like choosing a bandwidth: there is the usual tradeoff between bias and variance.

We use sieves informally all the time in the sense that we often fit more complex models when we have more data. The sieve idea was made formal by Grenander (1981) and Geman and Hwang (1982). Since then an enormous literature has developed. See Shen et al. (1999), Wong and Shen (1995), Shen and Wong (1994), Barron et al. (1999a), van de Geer (1995), Genovese and Wasserman (2000), and van de Geer (2000).

### 10.8 Shape-Restricted Inference

In the presence of shape restrictions, it is possible to make consistent, nonparametric inferences for a curve without imposing smoothness constraints. A typical example is estimating a regression function $r$ when $r$ is monotonic. A standard reference is Robertson et al. (1988).

Suppose that

$$
Y_{i}=r\left(x_{i}\right)+\epsilon_{i}, \quad i=1, \ldots, n
$$

where $x_{1}<\cdots<x_{n}, \mathbb{E}\left(\epsilon_{i}\right)=0$ and $\sigma^{2}=\mathbb{E}\left(\epsilon_{i}^{2}\right)$. Further, suppose that $r$ is nondecreasing. (This assumption can be tested as described in Section

[^25]10.9.) The least squares estimator $\widehat{r}_{n}$ is obtained by solving the restricted minimization:
$$
\operatorname{minimize} \sum_{i=1}^{n}\left(Y_{i}-r\left(x_{i}\right)\right)^{2} \quad \text { subject to } r \in \mathcal{F}_{\uparrow}
$$
where $\mathcal{F}_{\uparrow}$ is the set of nondecreasing functions. The resulting estimator $\widehat{r}_{n}$ is called the isotonic regression estimator.

The solution $\widehat{r}_{n}$ may be described as follows. Let $P_{0}=(0,0)$ and $P_{j}=$ $\left(j, \sum_{i=1}^{j} Y_{i}\right)$. Let $G(t)$ be the greatest convex minorant, that is, $G(t)$ is the supremum of all convex functions that lie below the points $P_{0}, \ldots, P_{n}$. Then $\widehat{r}_{n}$ is the left derivative of $G$.

The convex minorant $G$ can be found quickly using the pooled-adjacentviolators (PAV) algorithm. Start by joining all the points $P_{0}, P_{1}, \ldots$ with line segments. If the slope between $P_{0}$ and $P_{1}$ is greater than the slope between $P_{1}$ and $P_{2}$, replace these two segments with one line segment joining $P_{0}$ and $P_{2}$. If the slope between $P_{0}$ and $P_{2}$ is greater than the slope between $P_{2}$ and $P_{3}$, replace these two segments with one line segment joining $P_{0}$ and $P_{3}$. The process is continued in this way and the result is the function $G(t)$. See pages $8-10$ of Robertson et al. (1988) for more detail.

A number of results are available about the resulting estimator. For example, Zhang (2002) shows the following. If

$$
R_{n, p}(r)=\left(\frac{1}{n} \sum_{i=1}^{n} \mathbb{E}\left|\widehat{r}_{n}\left(x_{i}\right)-r\left(x_{i}\right)\right|^{p}\right)^{1 / p}
$$

where $1 \leq p \leq 3$, then

$$
\begin{equation*}
0.64+o(1) \leq \frac{n^{1 / 3}}{\sigma^{2 / 3} V^{1 / 3}} \sup _{V(r) \leq V} R_{n, p}(r) \leq M_{p}+o(1) \tag{10.30}
\end{equation*}
$$

where $V(r)$ is the total variation of $r$ and $M_{p}$ is a constant.
Optimal confidence bands are obtained in Dümbgen (2003), and Dümbgen and Johns (2004). Confidence bands for monotone densities are obtained in Hengartner and Stark (1995).

### 10.9 Testing

In this book we have concentrated on estimation and confidence sets. There is also a large literature on nonparametric testing. Many of the results can
be found in the monograph by Ingster and Suslina (2003). Other references include: Ingster (2002), Ingster (2001), Ingster and Suslina (2000), Ingster (1998), Ingster (1993a), Ingster (1993b), Ingster (1993c), Lepski and Spokoiny (1999), and Baraud (2002).

For example, let $Y_{i}=\theta_{i}+\epsilon_{i}$ where $\epsilon_{i} \sim N(0,1), i=1, \ldots, n$ and $\theta=$ $\left(\theta_{1}, \ldots, \theta_{n}\right)$ is unknown. Consider testing

$$
H_{0}: \theta=(0, \ldots, 0) \quad \text { versus } \quad H_{1}: \theta \in V_{n}=\left\{\theta \in \mathbb{R}^{n}:\|\theta\|_{p} \geq R_{n}\right\}
$$

where

$$
\|\theta\|_{p}=\left(\sum_{i=1}^{n}\left|\theta_{i}\right|^{p}\right)^{1 / p}
$$

for $0<p<\infty$. The type I and type II error of a test $\psi$ are

$$
\alpha_{n}(\psi)=\mathbb{E}_{0}(\psi), \quad \beta_{n}(\psi, \theta)=\mathbb{E}_{\theta}(1-\psi)
$$

Let

$$
\gamma_{n}=\inf _{\psi}\left(\alpha_{n}(\psi)+\sup _{\theta \in V_{n}} \beta_{n}(\psi, \theta)\right)
$$

be the smallest possible sum of the type I error and maximum type II error, over all tests. Ingster showed that

$$
\gamma_{n} \rightarrow 0 \Longleftrightarrow \frac{R_{n}}{R_{n}^{*}} \rightarrow \infty \quad \text { and } \quad \gamma_{n} \rightarrow 1 \Longleftrightarrow \frac{R_{n}}{R_{n}^{*}} \rightarrow 0
$$

where

$$
R_{n}^{*}=\left\{\begin{array}{cc}
n^{(1 / p)-(1 / 4)} & \text { if } p \leq 2 \\
n^{1 /(2 p)} & \text { if } p>2
\end{array}\right.
$$

Thus, $R_{n}^{*}$ is a critical rate that determines when the alternative is distinguishable. Results like these are intimately related to the results on confidence sets in Chapter 7.

The results for qualitative hypotheses are of a different nature. These are hypotheses like: $f$ is monotone, $f$ is positive, $f$ is convex and so on. The defining characteristic of such hypotheses is that they are cones, that is, they are closed under addition. For example, if $f$ and $g$ are monotone nondecreasing functions, then $f+g$ is also monotone nondecreasing. References include: Dümbgen and Spokoiny (2001), Baraud et al. (2003a), Baraud et al. (2003b), and Juditsky and Nemirovski (2002). Consider testing the null hypothesis that a regression function $r$ is nonincreasing. Suppose further that

$$
r \in\left\{f:[0,1] \rightarrow \mathbb{R}:|r(x)-r(y)| \leq L|x-y|^{s}, \text { for all } x, y \in[0,1]\right\}
$$

where $L>0$ and $0<s \leq 1$. Then there exist tests with uniformly large power for each function whose distance from the null is at least of order $L^{1 /(1+2 s)} n^{-s /(1+2 s)}$.

### 10.10 Computational Issues

I have completely neglected issues about efficient computation. Nonparametric methods work best with large data sets but implementing nonparametric methods with large data sets requires efficient computation.

Binning methods are quite popular for fast calculations. See Hall and Wand (1996), Fan and Marron (1994), Wand (1994), Holmström (2000), Sain (2002), and Scott (1992), for example. Chapter 12 of Loader (1999a) contains a good discussion of computation. In particular, there is a description of $\mathbf{k}$ d trees which are intelligently chosen partitions of the data that speed up computations. A number of publications on the use of k -d trees in statistics can be found at http://www. autonlab.org. Useful R code can be found at http://cran.r-project.org. The locfit program by Catherine Loader for local likelihood and local regression can be found at http://www.locfit.info.

### 10.11 Exercises

1. Consider the "errors in $Y$ " model:

$$
\begin{aligned}
Y_{i} & =r\left(X_{i}\right)+\epsilon_{i} \\
Y_{i}^{\bullet} & =Y_{i}+U_{i}, \quad \mathbb{E}\left(U_{i}\right)=0
\end{aligned}
$$

and the observed data are $\left(X_{1}, Y_{1}^{\bullet}\right), \ldots,\left(X_{n}, Y_{n}^{\bullet}\right)$. How does observing $Y_{i}^{\bullet}$ instead of $Y_{i}$ affect the estimation of $r(x)$ ?
2. Prove 10.1.
3. Prove equation (10.20).
4. Draw $X_{1}, \ldots, X_{n} \sim N(0,1)$ with $n=100$.
(a) Estimate the density using a kernel estimator.
(b) Let $W_{i}=X_{i}+\sigma_{u} U_{i}$ where $U_{i} \sim N(0,1)$. Compute the uncorrected and corrected density estimator from the $W_{i}$ s and compare the estimators. Try different values of $\sigma_{u}$.
(c) Repeat part (b) but let $U_{i}$ have a Cauchy distribution.
5. Generate 1000 observations from the model:

$$
\begin{aligned}
Y_{i} & =r\left(X_{i}\right)+\sigma_{\epsilon} \epsilon_{i} \\
W_{i} & =X_{i}+\sigma_{u} U_{i}
\end{aligned}
$$

where $r(x)=x+3 \exp \left(-16 x^{2}\right), \epsilon_{i} \sim N(0,1), U_{i} \sim N(0,1), X_{i} \sim$ $\operatorname{Unif}(-2,2), \sigma_{\epsilon}=0.5$, and $\sigma_{u}=0.1$.
(a) Use kernel regression to estimate $r$ from $\left(X_{1}, Y_{1}\right), \ldots,\left(X_{n}, Y_{n}\right)$. Find the bandwidth $h$ using cross-validation. Call the resulting estimator $r_{n}^{*}$. Use the bandwidth $h$ throughout the rest of the question.
(b) Use kernel regression to estimate $r$ from $\left(W_{1}, Y_{1}\right), \ldots,\left(W_{n}, Y_{n}\right)$.

Denote the resulting estimator by $\widehat{r}_{n}$.
(c) Compute the corrected estimator $\widetilde{r}_{n}$ given by (10.14).
(d) Compare $r, r_{n}^{*}, \widehat{r}_{n}$ and $\widetilde{r}_{n}$.
(e) Think of a way of finding a good bandwidth using only the $Y_{i}$ 's and the $W_{i}$ 's. Implement your method and compare the resulting estimator to the previous estimators.
6. Generate 1000 observations from the model:

$$
Y_{i}=\int K_{i}(s) d s+\sigma \epsilon_{i}
$$

where $r(x)=x+3 \exp \left(-16 x^{2}\right), \epsilon_{i} \sim N(0,1), \sigma=0.5, K_{i}(s)=$ $e^{-\left(s-x_{i}\right)^{2} / b^{2}}$, and $x_{i}=4(i / n)-2$. Try $b=.01, .1$ and 1 .
(a) Plot the Backus-Gilbert averaging kernel for several values of $x$. Interpret.
(b) Estimate $r$ using the method described in this chapter. Comment on the results.
7. Consider the infinite-dimensional Normal means model from Chapter 7 given by:

$$
Y_{i}=\theta_{i}+\frac{1}{\sqrt{n}} \epsilon_{i}, \quad i=1,2, \ldots
$$

Assume that $\theta=\left(\theta_{1}, \theta_{2}, \ldots\right)$ is in the Sobolev ellipsoid $\Theta=\{\theta$ : $\left.\sum_{j=1}^{\infty} \theta_{j}^{2} j^{2} \leq c^{2}\right\}$ for some $c>0$. Let the prior be such that each $\theta_{i}$ is independent and $\theta_{i} \sim N\left(0, \tau_{i}^{2}\right)$ for some $\tau_{i}>0$.
(a) Find the posterior of the $\theta_{i}$ 's. In particular, find the posterior mean $\widehat{\theta}$.
(b) Find conditions on the $\tau_{i}^{2}$ such that the posterior is consistent in the sense that

$$
\mathbb{P}_{\theta}(\|\theta-\widehat{\theta}\|>\epsilon) \rightarrow 0
$$

for any $\epsilon$, as $n \rightarrow \infty$.
(c) Let $\theta_{j}=1 / j^{4}$. Simulate from the model (take $n=100$ ), find the posterior and find $b_{n}$ such that $\mathbb{P}\left(||\widehat{\theta}-\theta|| \leq b_{n} \mid\right.$ Data $)=0.95$. Set $B_{n}=$ $\left\{\theta \in \Theta:\|\widehat{\theta}-\theta\| \leq b_{n}\right\}$. Repeat this whole process many times and estimate the frequentist coverage of $B_{n}$ (at this particular $\theta$ ) by counting how often $\theta \in B_{n}$. Report your findings.
(d) Repeat (c) for $\theta=(0,0, \ldots)$.

## Bibliography

Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. Second International Symposium on Information Theory 267-281.

Baraud, Y. (2002). Nonasymptotic minimax rates of testing in signal detection. Bernoulli 8 577-606.

Baraud, Y. (2004). Confidence balls in Gaussian regression. The Annals of Statistics 32 528-551.

Baraud, Y., Huet, S. and Laurent, B. (2003a). Adaptive tests of linear hypotheses by model selection. The Annals of Statistics 31 225-251.

Baraud, Y., Huet, S. and Laurent, B. (2003b). Adaptive tests of qualitative hypotheses. ESAIM PGS: Probability and Statistics 7 147-159.

Barron, A., Birge, L. and Massart, P. (1999a). Risk bounds for model selection via penalization. Probability Theory and Related Fields 113 301413.

Barron, A., Schervish, M. J. and Wasserman, L. (1999b). The consistency of posterior distributions in nonparametric problems. The Annals of Statistics 27 536-561.

Bellman, R. (1961). Adaptive Control Processes. Princeton University Press. Princeton, NJ.

Beran, R. (2000). REACT scatterplot smoothers: Superefficiency through basis economy. Journal of the American Statistical Association 95 155-171.

Beran, R. and Dümbgen, L. (1998). Modulation of estimators and confidence sets. The Annals of Statistics 26 1826-1856.

Bickel, P. J. and Freedman, D. A. (1981). Some asymptotic theory for the bootstrap. The Annals of Statistics 9 1196-1217.

Bickel, P. J., Klaassen, C. A. J., Ritov, Y. and Wellner, J. A. (1993). Efficient and adaptive estimation for semiparametric models. Johns Hopkins University Press. Baltimore, MD.

Blei, D., Griffiths, T., Jordan, M. and Tenebaum, J. (2004). Hierarchical topic models and the nested Chinese restaurant process. In $S$. Thrun, L. Saul, and B. Schoelkopf (Eds.), Advances in Neural Information Processing Systems (NIPS) 16, 2004. .

Blei, D. and Jordan, M. (2004). Variational methods for the dirichlet process. In Proceedings of the 21st International Conference on Machine Learning (ICML). Omnipress.

Breiman, L., Friedman, J. H., Olshen, R. A. and Stone, C. J. (1984). Classification and regression trees. Wadsworth. New York, NY.

Brown, L., Cai, T. and Zhou, H. (2005). A root-unroot transform and wavelet block thresholding approach to adaptive density estimation. unpublished.

Brown, L. D. and Low, M. G. (1996). Asymptotic equivalence of nonparametric regression and white noise. The Annals of Statistics 24 2384-2398.

Cai, T. and Low, M. (2005). Adaptive confidence balls. To appear: The Annals of Statistics .

Cai, T., Low, M. and Zhao, L. (2000). Sharp adaptive estimation by a blockwise method. Technical report, Wharton School, University of Pennsylvania, Philadelphia .

Carroll, R., Ruppert, D. and Stefanski, L. (1995). Measurement Error in Nonlinear Models. Chapman and Hall. New York, NY.

Casella, G. and Berger, R. L. (2002). Statistical Inference. Duxbury Press. New York, NY.

Cenc̆ov, N. (1962). Evaluation of an unknown distribution density from observations. Doklady 3 1559-1562.

Chaudhuri, P. and Marron, J. S. (1999). Sizer for exploration of structures in curves. Journal of the American Statistical Association 94 807-823.

Chaudhuri, P. and Marron, J. S. (2000). Scale space view of curve estimation. The Annals of Statistics 28 408-428.

Chen, S. S., Donoho, D. L. and Saunders, M. A. (1998). Atomic decomposition by basis pursuit. SIAM Journal on Scientific Computing 20 33-61.

Chen, S. X. and Qin, Y. S. (2000). Empirical likelihood confidence intervals for local linear smoothers. Biometrika 87 946-953.

Chen, S. X. and Qin, Y. S. (2002). Confidence intervals based on local linear smoother. Scandinavian Journal of Statistics 29 89-99.

Chu, C.-K. and Marron, J. S. (1991). Comparison of two bandwidth selectors with dependent errors. The Annals of Statistics 19 1906-1918.

Claeskens, G. and Hjort, N. (2004). Goodness-of-fit via nonparametric likelihood ratios. Scandinavian Journal of Statistics 31 487-513.

Cook, J. R. and Stefanski, L. A. (1994). Simulation-extrapolation estimation in parametric measurement error models. Journal of the American Statistical Association 89 1314-1328.

Cox, D. and Lewis, P. (1966). The Statistical Analysis of Series of Events. Chapman and Hall. New York, NY.

Cox, D. D. (1993). An analysis of Bayesian inference for nonparametric regression. The Annals of Statistics 21 903-923.

Cummins, D. J., Filloon, T. G. and Nychka, D. (2001). Confidence intervals for nonparametric curve estimates: Toward more uniform pointwise coverage. Journal of the American Statistical Association 96 233-246.

Daubechies, I. (1992). Ten Lectures on Wavelets. SIAM. New York, NY.

Davison, A. C. and Hinkley, D. V. (1997). Bootstrap Methods and Their Application. Cambridge University Press. Cambridge.

Devroye, L., Györfi, L. and Lugosi, G. (1996). A Probabilistic Theory of Pattern Recognition. Springer-Verlag. New York, NY.

Dey, D., Muller, P. and Sinha, D. (1998). Practical Nonparametric and Semiparametric Bayesian Statistics. Springer-Verlag. New York, NY.

Diaconis, P. and Freedman, D. (1986). On inconsistent Bayes estimates of location. The Annals of Statistics 14 68-87.

Donoho, D. L. and Johnstone, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. Biometrika 81 425-455.

Donoho, D. L. and Johnstone, I. M. (1995). Adapting to unknown smoothness via wavelet shrinkage. Journal of the American Statistical Association 90 1200-1224.

Donoho, D. L. and Johnstone, I. M. (1998). Minimax estimation via wavelet shrinkage. The Annals of Statistics 26 879-921.

Donoho, D. L., Johnstone, I. M., Kerkyacharian, G. and Picard, D. (1995). Wavelet shrinkage: Asymptopia? Journal of the Royal Statistical Society, Series B, Methodological 57 301-337.

Dümbgen, L. (2003). Optimal confidence bands for shape-restricted curves. Bernoulli 9 423-449.

Dümbgen, L. and Johns, R. (2004). Confidence bands for isotonic median curves using sign-tests. Journal of Computational and Graphical Statistics 13 519-533.

Dümbgen, L. and Spokoiny, V. G. (2001). Multiscale testing of qualitative hypotheses. The Annals of Statistics 29 124-152.

Efromovich, S. (1999). Nonparametric Curve Estimation: Methods, Theory and Applications. Springer-Verlag. New York, NY.

Efromovich, S. Y. and Pinsker, M. S. (1982). Estimation of squareintegrable probability density of a random variable. Problems of Information Transmission, (Transl of Problemy Peredachi Informatsii) 18 175-189.

Efromovich, S. Y. and Pinsker, M. S. (1984). A learning algorithm for nonparametric filtering. Automat. i Telemekh 11 58-65.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. The Annals of Statistics 7 1-26.

Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004). Least angle regression. The Annals of Statistics 32 407-499.

Efron, B. and Tibshirani, R. J. (1993). An Introduction to the Bootstrap. Chapman and Hall. New York, NY.

FAN, J. (1992). Design-adaptive nonparametric regression. Journal of the American Statistical Association 87 998-1004.

Fan, J. and Gijbels, I. (1995). Data-driven bandwidth selection in local polynomial fitting: Variable bandwidth and spatial adaptation. Journal of the Royal Statistical Society, Series B, Methodological 57 371-394.

Fan, J. and Gijbels, I. (1996). Local Polynomial Modelling and Its Applications. Chapman and Hall. New York, NY.

Fan, J. and Marron, J. S. (1994). Fast implementations of nonparametric curve estimators. Journal of Computational and Graphical Statistics 3 3556.

Fan, J. and Truong, Y. K. (1993). Nonparametric regression with errors in variables. The Annals of Statistics 21 1900-1925.

Faraway, J. J. (1990). Bootstrap selection of bandwidth and confidence bands for nonparametric regression. Journal of Statistical Computation and Simulation 37 37-44.

Faraway, J. J. and Sun, J. (1995). Simultaneous confidence bands for linear regression with heteroscedastic errors. Journal of the American Statistical Association 90 1094-1098.

Fernholz, L. T. (1983). Von Mises' Calculus for Statistical Functionals. Springer-Verlag. New York, NY.

Freedman, D. (1999). Wald lecture: On the Bernstein-von Mises theorem with infinite-dimensional parameters. The Annals of Statistics 271119 1141.

Friedman, J. H. (1991). Multivariate adaptive regression splines. The Annals of Statistics 19 1-67.

Friedman, J. H. and Stuetzle, W. (1981). Projection pursuit regression. Journal of the American Statistical Association 76 817-823.

Fuller, W. A. (1987). Measurement Error Models. John Wiley. New York, NY.

Gao, F., Wahba, G., Klein, R. and Klein, B. (2001). Smoothing spline ANOVA for multivariate Bernoulli observations with application to ophthalmology data. Journal of the American Statistical Association 96 127-160.

Gasser, T., Sroka, L. and Jennen-Steinmetz, C. (1986). Residual variance and residual pattern in nonlinear regression. Biometrika 73 625-633.

Geman, S. and Hwang, C.-R. (1982). Nonparametric maximum likelihood estimation by the method of sieves. The Annals of Statistics 10 401-414.

Genovese, C., Miller, C., Nichol, R., Arjunwadkar, M. and Wasserman, L. (2004). Nonparametric inference for the cosmic microwave background. Statistical Science 19 308-321.

Genovese, C. and Wasserman, L. (2005). Nonparametric confidence sets for wavelet regression. Annals of Statistics 33 698-729.

Genovese, C. R. and Wasserman, L. (2000). Rates of convergence for the Gaussian mixture sieve. The Annals of Statistics 28 1105-1127.

Ghosal, S., Ghosh, J. K. and Ramamoorthi, R. V. (1999). Posterior consistency of Dirichlet mixtures in density estimation. The Annals of Statistics 27 143-158.

Ghosal, S., Ghosh, J. K. and Van Der Vaart, A. W. (2000). Convergence rates of posterior distributions. The Annals of Statistics 28 500-531.

Ghosh, J. and Ramamoorthi, R. (2003). Bayesian Nonparametrics. Springer-Verlag. New York, NY.

Glad, I., Hjort, N. and Ushakov, N. (2003). Correction of density estimators that are not densities. Scandinavian Journal of Statististics 30 415-427.

Goldenshluger, A. and Nemirovski, A. (1997). On spatially adaptive estimation of nonparametric regression. Mathematical Methods of Statistics 6 135-170.

Green, P. J. and Silverman, B. W. (1994). Nonparametric regression and generalized linear models: a roughness penalty approach. Chapman and Hall. New York, NY.

Grenander, U. (1981). Abstract Inference. John Wiley. New York, NY.
Hall, P. (1987). On Kullback-Leibler loss and density estimation. The Annals of Statistics 15 1491-1519.

Hall, P. (1992a). The Bootstrap and Edgeworth Expansion. Springer-Verlag. New York, NY.

Hall, P. (1992b). On bootstrap confidence intervals in nonparametric regression. The Annals of Statistics 20 695-711.

Hall, P. (1993). On Edgeworth expansion and bootstrap confidence bands in nonparametric curve estimation. Journal of the Royal Statistical Society, Series B, Methodological 55 291-304.

Hall, P. and Wand, M. P. (1996). On the accuracy of binned kernel density estimators. Journal of Multivariate Analysis 56 165-184.

Härdle, W. (1990). Applied Nonparametric Regression. Cambridge University Press. Cambridge.

Härdle, W. and Bowman, A. W. (1988). Bootstrapping in nonparametric regression: Local adaptive smoothing and confidence bands. Journal of the American Statistical Association 83 102-110.

Härdle, W., Hall, P. and Marron, J. S. (1988). How far are automatically chosen regression smoothing parameters from their optimum? Journal of the American Statistical Association 83 86-95.

Härdle, W., Kerkyacharian, G., Picard, D. and Tsybakov, A. (1998). Wavelets, Approximation, and Statistical Applications. Springer-Verlag. New York, NY.

Härdle, W. and Mammen, E. (1993). Comparing nonparametric versus parametric regression fits. The Annals of Statistics 21 1926-1947.

Härdle, W. and Marron, J. S. (1991). Bootstrap simultaneous error bars for nonparametric regression. The Annals of Statistics 19 778-796.

Hastie, T. and Loader, C. (1993). Local regression: Automatic kernel carpentry. Statistical Science 8 120-129.

Hastie, T. and Tibshirani, R. (1999). Generalized Additive Models. Chapman and Hall. New York, NY.

Hastie, T., Tibshirani, R. and Friedman, J. H. (2001). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. SpringerVerlag. New York, NY.

Hengartner, N. W. and Stark, P. B. (1995). Finite-sample confidence envelopes for shape-restricted densities. The Annals of Statistics 23 525550.

Hjort, N. (1999). Towards semiparametric bandwidth selectors for kernel density estimation. Statistical Research Report. Department of Mathematics, University of Oslo .

Hjort, N. (2003). Topics in nonparametric Bayesian statistics. In Highly Structured Stochastic Systems. P. Green, N.L. Hjort, S. Richardson (Eds.). Oxford University Press. Oxford.

Hjort, N. L. and Jones, M. C. (1996). Locally parametric nonparametric density estimation. The Annals of Statistics 24 1619-1647.

Holmström, L. (2000). The accuracy and the computational complexity of a multivariate binned kernel density estimator. Journal of Multivariate Analysis 72 264-309.

Hotelling, H. (1939). Tubes and spheres in $n$-spaces, and a class of statistical problems. American Journal of Mathematics 61 440-460.

Huang, T.-M. (2004). Convergence rates for posterior distributions and adaptive estimation. The Annals of Statistics 32 1556-1593.

Ibragimov, I. A. and Has'minskil, R. Z. (1977). On the estimation of an infinite-dimensional parameter in Gaussian white noise. Soviet Math. Dokl. 236 1053-1055.

Ingster, Y. and Suslina, I. (2003). Nonparametric Goodness-of-Fit Testing Under Gaussian Models. Springer-Verlag. New York, NY.

Ingster, Y. I. (1993a). Asymptotically minimax hypothesis testing for nonparametric alternatives. I. Mathematical Methods of Statistics 2 85-114.

Ingster, Y. I. (1993b). Asymptotically minimax hypothesis testing for nonparametric alternatives. II. Mathematical Methods of Statistics 2 171-189.

Ingster, Y. I. (1993c). Asymptotically minimax hypothesis testing for nonparametric alternatives, III. Mathematical Methods of Statistics 2 249-268.

Ingster, Y. I. (1998). Minimax detection of a signal for $l^{n}$-balls. Mathematical Methods of Statistics 7 401-428.

Ingster, Y. I. (2001). Adaptive detection of a signal of growing dimension. I. Mathematical Methods of Statistics 10 395-421.

Ingster, Y. I. (2002). Adaptive detection of a signal of growing dimension. II. Mathematical Methods of Statistics 11 37-68.

Ingster, Y. I. and Suslina, I. A. (2000). Minimax nonparametric hypothesis testing for ellipsoids and Besov bodies. ESAIM PGSS: Probability and Statistics 4 53-135.

Jang, W., Genovese, C. and Wasserman, L. (2004). Nonparametric confidence sets for densities. Technical Report, Carnegie Mellon University, Pittsburgh .

Johnstone, I. (2003). Function Estimation in Gaussian Noise: Sequence Models. Unpublished manuscript.

Juditsky, A. and Lambert-Lacroix, S. (2003). Nonparametric confidence set estimation. Mathematical Methods of Statistics 19 410-428.

Juditsky, A. and Nemirovski, A. (2002). On nonparametric tests of positivity/monotonicity/convexity. The Annals of Statistics 30 498-527.

Lepski, O. (1999). How to improve the accuracy of estimation. Mathematical Methods in Statistics 8 441-486.

Lepski, O. V., Mammen, E. and Spokoiny, V. G. (1997). Optimal spatial adaptation to inhomogeneous smoothness: An approach based on kernel estimates with variable bandwidth selectors. The Annals of Statistics $\mathbf{2 5}$ 929-947.

Lepski, O. V. and Spokoiny, V. G. (1999). Minimax nonparametric hypothesis testing: The case of an inhomogeneous alternative. Bernoulli 5 333-358.

Lepskir, O. V. (1991). On a problem of adaptive estimation in Gaussian white noise. Theory of Probability and Its Applications (Transl of Teorija Verojatnostei i ee Primenenija) 35 454-466.

Li, K.-C. (1989). Honest confidence regions for nonparametric regression. The Annals of Statistics 17 1001-1008.

Lin, X., Wahba, G., Xiang, D., Gao, F., Klein, R. and Klein, B. (2000). Smoothing spline ANOVA models for large data sets with Bernoulli observations and the randomized GACV. The Annals of Statistics 28 1570-1600.

Lin, Y. (2000). Tensor product space ANOVA models. The Annals of Statistics 28 734-755.

Loader, C. (1999a). Local Regression and Likelihood. Springer-Verlag. New York, NY.

Loader, C. R. (1999b). Bandwidth selection: classical or plug-in? The Annals of Statistics 27 415-438.

Low, M. G. (1997). On nonparametric confidence intervals. The Annals of Statistics 25 2547-2554.

Mallows, C. L. (1973). Some comments on $C_{p}$. Technometrics 15 661-675.
Mammen, E. and van de Geer, S. (1997). Locally adaptive regression splines. The Annals of Statistics 25 387-413.

Marron, J. S. and Wand, M. P. (1992). Exact mean integrated squared error. The Annals of Statistics 20 712-736.

McAuliffe, J., Blei, D. and Jordan, M. (2004). Variational inference for Dirichlet process mixtures. Department of Statistics, University of California, Berkeley.

McCullagh, P. and Nelder, J. A. (1999). Generalized linear models. Chapman and Hall. New York, NY.

Morgan, J. N. and Sonquist, J. A. (1963). Problems in the analysis of survey data, and a proposal. Journal of the American Statistical Association 58 415-434.

Müller, H.-G. and Stadtmller, U. (1987). Variable bandwidth kernel estimators of regression curves. The Annals of Statistics 15 182-201.

Naiman, D. Q. (1990). Volumes of tubular neighborhoods of spherical polyhedra and statistical inference. The Annals of Statistics 18 685-716.

Neumann, M. H. (1995). Automatic bandwidth choice and confidence intervals in nonparametric regression. The Annals of Statistics 23 1937-1959.

Neumann, M. H. and Polzehl, J. (1998). Simultaneous bootstrap confidence bands in nonparametric regression. Journal of Nonparametric Statistics 9 307-333.

Nussbaum, M. (1985). Spline smoothing in regression models and asymptotic efficiency in $L_{2}$. The Annals of Statistics 13 984-997.

Nussbaum, M. (1996a). Asymptotic equivalence of density estimation and Gaussian white noise. The Annals of Statistics 24 2399-2430.

Nussbaum, M. (1996b). The Pinsker bound: A review. In Encyclopedia of Statistical Sciences (S. Kotz, Ed). Wiley. New York, NY.

Nychka, D. (1988). Bayesian confidence intervals for smoothing splines. Journal of the American Statistical Association 83 1134-1143.

Ogden, R. T. (1997). Essential Wavelets for Statistical Applications and Data Analysis. Birkhäuser. Boston, MA.

Opsomer, J., Wang, Y. and Yang, Y. (2001). Nonparametric regression with correlated errors. Statistical Science 16 134-153.

O'Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. Statistical Science 1 502-527.

Pagano, M. and Gauvreau, K. (1993). Principles of biostatistics. Duxbury Press. New York, NY.

Parzen, E. (1962). On estimation of a probability density function and mode. The Annals of Mathematical Statistics 33 1065-1076.

Picard, D. and Tribouley, K. (2000). Adaptive confidence interval for pointwise curve estimation. The Annals of Statistics 28 298-335.

Quenouille, M. (1949). Approximate tests of correlation in time series. Journal of the Royal Statistical Society B11 18-84.

Rice, J. (1984). Bandwidth choice for nonparametric regression. The Annals of Statistics 12 1215-1230.

Rice, S. (1939). The distribution of the maxima of a random curve. American Journal of Mathematics 61 409-416.

Robertson, T., Wright, F. T. and Dykstra, R. (1988). Order restricted statistical inference. Wiley. New York, NY.

Robins, J. and van der Vaart, A. (2005). Adaptive nonparametric confidence sets. To appear: The Annals of Statistics .

Rosenblatt, M. (1956). Remarks on some nonparametric estimates of a density function. Annals of Mathematical Statistics 27 832-837.

Rudemo, M. (1982). Empirical choice of histograms and kernel density estimators. Scandinavian Journal of Statistics 965-78.

Ruppert, D. (1997). Empirical-bias bandwidths for local polynomial nonparametric regression and density estimation. Journal of the American Statistical Association 92 1049-1062.

Ruppert, D., Wand, M. and Carroll, R. (2003). Semiparametric Regression. Cambridge University Press. Cambridge.

Ruppert, D. and Wand, M. P. (1994). Multivariate locally weighted least squares regression. The Annals of Statistics 22 1346-1370.

Sain, S. R. (2002). Multivariate locally adaptive density estimation. Computational Statistics and Data Analysis 39 165-186.

Scheffé, H. (1959). The Analysis of Variance. Wiley. New York, NY.
Schwartz, L. (1965). On Bayes procedures. Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 4 10-26.

Schwarz, G. (1978). Estimating the dimension of a model. The Annals of Statistics 6 461-464.

Scott, D. W. (1992). Multivariate Density Estimation: Theory, Practice, and Visualization. Wiley. New York, NY.

Serfling, R. J. (1980). Approximation Theorems of Mathematical Statistics. Wiley. New York, NY.

Shao, J. and Tu, D. (1995). The Jackknife and Bootstrap. Springer-Verlag. New York, NY.

Shen, X., Shi, J. and Wong, W. H. (1999). Random sieve likelihood and general regression models. Journal of the American Statistical Association 94 835-846.

Shen, X. and Wasserman, L. (2001). Rates of convergence of posterior distributions. The Annals of Statistics 29 687-714.

Shen, X. and Wong, W. H. (1994). Convergence rate of sieve estimates. The Annals of Statistics 22 580-615.

Sigrist, M. E. (1994). Air Monitoring by Spectroscopic Techniques. Wiley. New York, NY.

Silverman, B. W. (1984). Spline smoothing: The equivalent variable kernel method. The Annals of Statistics 12 898-916.

Silverman, B. W. (1986). Density Estimation for Statistics and Data Analysis. Chapman and Hall. New York, NY.

Simonoff, J. S. (1996). Smoothing Methods in Statistics. Springer-Verlag. New York, NY.

Singh, K. (1981). On the asymptotic accuracy of Efron's bootstrap. The Annals of Statistics 9 1187-1195.

Staudenmayer, J. and Ruppert, D. (2004). Local polynomial regression and simulation-extrapolation. Journal of the Royal Statistical Society Series $B 66$ 17-30.

Stefanski, L. A. (1985). The effects of measurement error on parameter estimation. Biometrika 72 583-592.

Stefanski, L. A. (1990). Rates of convergence of some estimators in a class of deconvolution problems. Statistics and Probability Letters 9 229-235.

Stefanski, L. A. and Cook, J. R. (1995). Simulation-extrapolation: The measurement error jackknife. Journal of the American Statistical Association 90 1247-1256.

Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. The Annals of Statistics 9 1135-1151.

Stone, C. J. (1984). An asymptotically optimal window selection rule for kernel density estimates. The Annals of Statistics 12 1285-1297.

Sun, J. and Loader, C. R. (1994). Simultaneous confidence bands for linear regression and smoothing. The Annals of Statistics 22 1328-1345.

Teh, Y., Jordan, M., Beal, M. and Blei, D. (2004). Hierarchical dirichlet processes. In Technical Report. Department of Statistics, University of California, Berkeley.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society, Series B, Methodological 58 267288.

Tukey, J. (1958). Bias and confidence in not quite large samples. The Annals of Mathematical Statistics 29614.
van de Geer, S. (1995). The method of sieves and minimum contrast estimators. Mathematical Methods of Statistics 420-38.
van de Geer, S. A. (2000). Empirical Processes in M-Estimation. Cambridge University Press. Cambridge.
van der Vaart, A. W. (1998). Asymptotic Statistics. Cambridge University Press.
van der Vaart, A. W. and Wellner, J. A. (1996). Weak Convergence and Empirical Processes: With Applications to Statistics. Springer-Verlag.

Venables, W. N. and Ripley, B. D. (2002). Modern Applied Statistics with $S$. Springer-Verlag. New York, NY.

Wahba, G. (1983). Bayesian "confidence intervals" for the cross-validated smoothing spline. Journal of the Royal Statistical Society, Series B, Methodological 45 133-150.

Wahba, G. (1990). Spline models for observational data. SIAM. New York, NY.

Wahba, G., Wang, Y., Gu, C., Klein, R. and Klein, B. (1995). Smoothing spline ANOVA for exponential families, with application to the Wisconsin Epidemiological Study of Diabetic Retinopathy. The Annals of Statistics 23 1865-1895.

Walker, S. (2004). Modern Bayesian asymptotics. Statistical Science 19 111-117.

Walker, S. and Hjort, N. L. (2001). On Bayesian consistency. Journal of the Royal Statistical Society, Series B, Methodological 63 811-821.

Wand, M. P. (1994). Fast computation of multivariate kernel estimators. Journal of Computational and Graphical Statistics 3 433-445.

Wasserman, L. (1998). Asymptotic properties of nonparametric Bayesian procedures. In Practical Nonparametric and Semiparametric Bayesian Statistics. Springer-Verlag. New York, NY.

Wasserman, L. (2000). Bayesian model selection and model averaging. Journal of Mathematical Psychology 44 92-107.

Wasserman, L. (2004). All of Statistics: A Concise Course in Statistical Inference. Springer-Verlag. New York, NY.

Weisberg, S. (1985). Applied Linear Regression. Wiley. New York, NY.
Wong, W. H. and Shen, X. (1995). Probability inequalities for likelihood ratios and convergence rates of sieve MLEs. The Annals of Statistics 23 339-362.

Yu, K. and Jones, M. (2004). Likelihood-based local linear estimation of the conditional variance function. Journal of the American Statistical Association 99 139-144.

Zhang, C.-H. (2002). Risk bounds in isotonic regression. The Annals of Statistics 2528 - 555.

Zhang, P. (1991). Variable selection in nonparametric regression with continuous covariates. The Annals of Statistics 19 1869-1882.

Zhao, L. H. (2000). Bayesian aspects of some nonparametric problems. The Annals of Statistics 28 532-552.

## List of Symbols

| $\mathbb{R}$ | real numbers |
| :--- | :--- |
| $\mathbb{P}(A)$ | probability of event $A$ |
| $F_{X}$ | cumulative distribution function |
| $f_{X}$ | probability density (or mass) function |
| $X \sim F$ | $X$ has distribution $F$ |
| $X \sim f$ | $X$ has distribution with density $f$ |
| $X \stackrel{\text { d }}{=} Y$ | $X$ and $Y$ have the same distribution |
| IID | independent and identically distributed |
| $X_{1}, \ldots, X_{n} \sim F$ | IID sample of size $n$ from $F$ |
| $\phi$ | standard Normal probability density |
| $\Phi$ | standard Normal distribution function |
| $z_{\alpha}$ | upper $\alpha$ quantile of $N(0,1): z_{\alpha}=\Phi^{-1}(1-\alpha)$ |
| $\mathbb{E}(X)=\int x d F(x)$ | expected value (mean) of random variable $X$ |
| $\mathbb{V}(X)$ | variance of random variable $X$ |
| $\operatorname{Cov}(X, Y)$ | covariance between $X$ and $Y$ |
| $\xrightarrow{\mathrm{P}}$ | convergence in probability |
| $\rightsquigarrow$ | convergence in distribution |
| $\longrightarrow$ | $x_{n} / a_{n} \rightarrow 0$ |
| $x_{n}=o\left(a_{n}\right)$ | $\left\|x_{n} / a_{n}\right\|$ is bounded for large $n$ |
| $x_{n}=O\left(a_{n}\right)$ | $X_{n} / a_{n} \xrightarrow{\mathrm{P}} 0$ |
| $X_{n}=o_{P}\left(a_{n}\right)$ | $\left\|X_{n} / a_{n}\right\|$ is bounded in probability for large $n$ |
| $X_{n}=O_{P}\left(a_{n}\right)$ | statistical functional (the mean, for example) |
| $T(F)$ | likelihood function |


| Distribution | PDF or probability function | mean | variance | MGF |
| :---: | :---: | :---: | :---: | :---: |
| Point mass at $a$ | $I(x=a)$ | $a$ | 0 | $e^{a t}$ |
| Bernoulli $(p)$ | $p^{x}(1-p)^{1-x}$ | $p$ | $p(1-p)$ | $p e^{t}+(1-p)$ |
| $\operatorname{Binomial}(n, p)$ | $\binom{n}{x} p^{x}(1-p)^{n-x}$ | $n p$ | $n p(1-p)$ | $\left(p e^{t}+(1-p)\right)^{n}$ |
| Geometric $(p)$ | $p(1-p)^{x-1} I(x \geq 1)$ | $1 / p$ | $\frac{1-p}{p^{2}}$ | $\frac{p e^{t}}{1-(1-p) e^{t}}(t<-\log (1-p))$ |
| $\operatorname{Poisson}(\lambda)$ | $\frac{\lambda^{x} e^{-\lambda}}{x!}$ | $\lambda$ | $\lambda$ | $e^{\lambda\left(e^{t}-1\right)}$ |
| Uniform $(a, b)$ | $I(a<x<b) /(b-a)$ | $\frac{a+b}{2}$ | $\frac{(b-a)^{2}}{12}$ | $\frac{e^{b t}-e^{a t}}{(b-a) t}$ |
| $\operatorname{Normal}\left(\mu, \sigma^{2}\right)$ | $\frac{1}{\sigma \sqrt{2 \pi}} e^{-(x-\mu)^{2} /\left(2 \sigma^{2}\right)}$ | $\mu$ | $\sigma^{2}$ | $\exp \left\{\mu t+\frac{\sigma^{2} t^{2}}{2}\right\}$ |
| Exponential $(\beta)$ | $\frac{e^{-x / \beta}}{\beta}$ | $\beta$ | $\beta^{2}$ | $\frac{1}{1-\beta t}(t<1 / \beta)$ |
| $\operatorname{Gamma}(\alpha, \beta)$ | $\frac{x^{\alpha-1} e^{-x / \beta}}{\Gamma(\alpha) \beta^{\alpha}}$ | $\alpha \beta$ | $\alpha \beta^{2}$ | $\left(\frac{\beta}{\beta-t}\right)^{\alpha}(t<\beta)$ |
| $\operatorname{Beta}(\alpha, \beta)$ | $\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha) \Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$ | $\frac{\alpha}{\alpha+\beta}$ | $\frac{\alpha \beta}{(\alpha+\beta)^{2}(\alpha+\beta+1)}$ | $1+\sum_{k=1}^{\infty}\left(\prod_{r=0}^{k-1} \frac{\alpha+r}{\alpha+\beta+r}\right) \frac{t^{k}}{k!}$ |
| $t_{\nu}$ | $\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)} \frac{1}{\left(1+\frac{x^{2}}{\nu}\right)^{(\nu+1) / 2}} .$ | $0($ if $\nu>1)$ | $\frac{\nu}{\nu-2}($ if $\nu>2)$ | does not exist |
| $\chi_{p}^{2}$ | $\frac{1}{\Gamma(p / 2) 2^{p / 2}} x^{(p / 2)-1} e^{-x / 2}$ | $p$ | $2 p$ | $\left(\frac{1}{1-2 t}\right)^{p / 2}(t<1 / 2)$ |

## Index

$C_{p}$ statistic, 70
$L_{p}$ loss, 57
$\chi^{2}$ confidence set, 159
$\sigma$-field, 3
$\mathrm{BC}_{a}$ intervals, 34
REACT, 183, 188
SIMEX, 229
SURE, 150
adaptive estimation, 158
adaptive inference, 115
adaptive kernels, 136
additive model, 50, 103
adjusted percentile methods, 34
AIC, 181
all possible subsets, 153
almost sure convergence, 4
attenuation bias, 228
average coverage, 94
average mean squared error, 52
averaging kernel, 234
backfitting, 103
Backus-Gilbert averaging kernel, 234
bandwidth, 71, 132
Baraud's confidence set, 161
Bart Simpson, 125
basis, 147
basis pursuit, 171, 216
Bayes rule, 237
Beran-Dümbgen-Stein pivotal method, 165

Bernstein's inequality, 9
Berry-Esséen bound, 9
Besov seminorm, 213
Besov space, 149, 211, 213
bias problem, 89
bias-variance decomposition, 5
bias-variance tradeoff, 50
bias-corrected and accelerated, 34
BIC, 181
Big Bang, 45
binning, 140
bins, $\underline{127}$
binwidth, $\underline{127}$
Bonferroni inequality, 94
bootstrap, 30
parametric, 31
bootstrap confidence interval, 32
percentile, 34
pivotal, 32
studentized pivotal, 33
boundary bias, 73
boxcar kernel, 55
cascade equations, 222
Cauchy-Schwartz inequality, 9
central limit theorem, 4
Chebyshev's inequality, 8
classification, 236
claw, 125
CMB, 46
column space, 63
complete, 147
confidence ball, 6
confidence band, $6,15,64$
confidence envelope, 6
confidence interval, 6
bootstrap, 32
pivotal, 32
confidence set, 5, 57, 159
$\chi^{2}, 159$
Baraud, 161
Beran-Dümbgen-Stein method, 165
finite sample, 6
Lepski, 161
pointwise asymptotic, 6
uniform asymptotic, 6
continuous mapping theorem, 4
convergence in distribution, 3
convergence in probability, 3
convex function, 9
correlation, 17
cosine basis, 148
cosmic microwave background, 46
covariate, 61
cross-validation, 126, 127, 135
cubic spline, 81
curse of dimensionality, $58,100,138$, 193
deconvolution, 232
delta method, 4, 7
density estimation, 43, 125
orthogonal series, 192
design bias, 73
detail coefficients, 200
dictionary, 215
dilation equation, 204
Doppler function, 77,188
double bootstrap, 35
Dvoretzky-Kiefer-Wolfowitz (DKW) inequality, 14
effective degrees of freedom, 64,67
ellipsoid, 148, 154
empirical detail coefficients, 207
empirical distribution function, 2 , 13
empirical influence function, 18
empirical probability distribution, 21
empirical scaling coefficients, 207
Epanechnikov kernel, 55
errors-in-variables, 116, 227
estimated risk, 127
events, 3
fast Fourier transform, 136
father wavelet, 203
feature, 61
filter, 223
Fisher information, 5
fitted values, 66
Fourier basis, 149
Fourier transform, 199

Gateâux derivative, 18, 23
Gaussian kernel, 55
Gaussian process, 92
generalized cross-validation, 70
generalized linear model, 66, 96
Glivenko-Cantelli theorem, 14
greatest convex minorant, 238
greedy modulator, 195

Hölder space, 213
Haar father wavelet, 199
Haar scaling function, 199
Haar wavelet, 199
Hadamard differentiable, 19, 23
hard thresholding, 152, 208
hat matrix, 66
Heisenberg uncertainty principle, 222
heteroscedastic, 87
histogram, 127
Hoeffding's inequality, 8
homoscedasticity, 87
Huber's function, 116

ICI, 216, 225
ideal linear estimator, 155
idempotent, 64
ill-posed, 235
independent component analysis, 117
influence function, 18, 29
inner product, 147
integrated mean squared error, 52 , 126
integrated risk, 52,174
integrated squared error, 126
interquartile range, 135
intersecting confidence intervals, 216
inverse Fourier transform, 199
inverse problems, 233
irregular design, 190
isotonic regression, 238
jackknife, 27
jackknife bias estimate, 27
James-Stein estimator, 153, 155, 185
minimaxity, 157
Jensen's inequality, 9
k-d trees, 240
kernel, 55, 131
kernel density estimator, 131, 132
kernel estimator, 71
kernel regression, 71
kill it or keep it, 210
knots, 81
Kullback-Leibler loss, 57

LARS, 216
lasso, 171, 216
leave-one-out cross-validation, 69, 126
Legendre basis, 148
Lepski's confidence set, 161
likelihood function, 5
linear functional, 15
linear minimax risk, 113
linear shrinkage estimators, 155
linear smoother, 66
Lipschitz, 122
local linear regression, 76
local log-likelihood, 97
local polynomial, 75
local polynomial log-likelihood, 138
local polynomial regression, 74
locally adaptive, 115,197
log-likelihood function, 5
logistic regression, 48
logistic regression model, 65
loss function, 57
Mallows' $C_{p}, 181$
Mallows' metric, 37
Markov's inequality, 8
MARS, 110
maximum likelihood estimator, 5
mean squared error, 44, 51, 68
measurement error, 116, 227
method of moments estimator, 229
Mill's inequality, 8
minimax risk, 113,153
minimax theory, 145
mirror, 215
model selection, 152
modified cross-validation, 236
modulator, 185
monotone modulator, 185
mother Haar wavelet, 199
mother wavelet, 205
MRA, 203
MSE, 51, 68
multiresolution analysis, 203
multivariate adaptive regression splines, 110

Nadaraya-Watson kernel estimator, $\underline{71}$
nested subset selection, 153, 185
nonparametric delta method, 20
nonparametric regression, 61, 71
norm, 147
normal means model, 145
normal reference rule, 135
normal-based interval, 16
olive nonparametric statistics, i
oracle, $155,157,171,172,210$
orthonormal, 147
outlier resistant, 57
oversmoothing, 44
parametric bootstrap, 31
partitioned cross-validation, 236
PAV, 188, 238
penalization, 171
penalized regression, 81
penalized sums of squares, 81
percentile interval, 34
periodic Sobolev class, 148
pilot estimate, 112
Pinsker constant, 154
Pinsker's theorem, 153
pivot, 165
pivotal interval, 32
studentized, 33
plug-in bandwidth, 111, 136
plug-in estimator, 15
pooled-adjacent-violators, 188, 238
predictive risk, 53
principal component analysis, 117
probability measure, 3
probability space, 3
projection pursuit, 117
projection pursuit regression, 104
pseudo-values, 28
quadratic form, 86
quantile regression, 116
quantiles, 17
random variable, 3
react, 183
redshift, 44
regression
kernel, 71
nonparametric, 71
regression splines, 85
regression tree, 108
regressogram, 48, 67
regular design, 184
regularization, 81
residuals, 64
response variable, 61
reweighted least squares, 65
ridge regression, 84
risk, 126
risk estimation and adaptation after coordinate transformation, 183
robust regression, 116
roughness penalty, 81
sample quantile, 17
sample space, 3
scale-space, 140
scale-space smoothing, 99
scaling coefficients, 200, 204
scaling function, 203
score function, 5
semiparametric models, 235
shatter, 22
shrinkage, 155
sieve, 237
signal, 223
simulation extrapolation, 229
skewness, 17
Slutsky's theorem, 4
smoother, 61
smoothing, 43
smoothing bias, 89
smoothing matrix, 66
smoothing parameter, 44
Sobolev space, 148
soft threshold estimator, 152
soft thresholding, 208
sparse, 199, 208
sparsity, 171
spatially adaptive, 115,197
spatially inhomogeneous, 77,115 , 197
splines, 81
squared error, 50
squared prediction error, 53
statistical functional, 15
statistical model, 2
Stein's lemma, 151
Stein's unbiased risk estimator, 150
Stone's theorem, 136
strawberry fields forever, 7
strong law of large numbers, 4
studentized pivotal interval, 33
SureShrink, 210
surgery, 192
symmetric, 64
symmlet, 205
tensor product, 193
thin-plate spline, 102
trace, 64
training error, 68
tricube kernel, 55
truncated power basis, 82
tube formula, 92, 119
two-scale relationship, 204
undersmoothing, 44, 90, 117
upcrossing, 120

Vapnik-Chervonenkis (VC) theory, 22
variable bandwidth selection, 115
variable selection, 117
VC dimension, 22
VisuShrink, 210
wavelets, 197
weak convergence, 3
weak differentiability, 173
weak law of large numbers, 4
white noise model, 173
wiggly, 2, 115

Madansky: Prescriptions for Working Statisticians
McPherson: Applying and Interpreting Statistics: A Comprehensive Guide, Second Edition
Mueller: Basic Principles of Structural Equation Modeling: An Introduction to LISREL and EQS
Nguyen and Rogers: Fundamentals of Mathematical Statistics: Volume I: Probability for Statistics
Nguyen and Rogers: Fundamentals of Mathematical Statistics: Volume II: Statistical Inference
Noether: Introduction to Statistics: The Nonparametric Way
Nolan and Speed: Stat Labs: Mathematical Statistics Through Applications
Peters: Counting for Something: Statistical Principles and Personalities
Pfeiffer: Probability for Applications
Pitman: Probability
Rawlings, Pantula and Dickey: Applied Regression Analysis
Robert: The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation, Second Edition
Robert and Casella: Monte Carlo Statistical Methods
Rose and Smith: Mathematical Statistics with Mathematica
Santner and Duffy: The Statistical Analysis of Discrete Data
Saville and Wood: Statistical Methods: The Geometric Approach
Sen and Srivastava: Regression Analysis: Theory, Methods, and Applications
Shao: Mathematical Statistics, Second Edition
Shorack: Probability for Statisticians
Shumway and Stoffer: Time Series Analysis and Its Applications
Simonoff: Analyzing Categorical Data
Terrell: Mathematical Statistics: A Unified Introduction
Timm: Applied Multivariate Analysis
Toutenburg: Statistical Analysis of Designed Experiments, Second Edition
Wasserman: All of Nonparametric Statistics
Wasserman: All of Statistics: A Concise Course in Statistical Inference
Whittle: Probability via Expectation, Fourth Edition
Zacks: Introduction to Reliability Analysis: Probability Models and Statistical Methods


[^0]:    ${ }^{1}$ There exist tighter confidence bands but we use the DKW band because it is simple.

[^1]:    ${ }^{2}$ Hadamard differentiability is defined in the appendix.

[^2]:    ${ }^{1}$ When an object moves away from us, its light gets shifted towards the red end of the spectrum, called redshift. The faster an object moves away from us, the more its light is redshifted. Objects further away are moving away from us faster than close objects. Hence the redshift can be used to deduce the distance. This is more complicated than it sounds because the conversion from redshift to distance requires knowledge of the geometry of the universe.

[^3]:    ${ }^{1}$ If you want to peek ahead, Figure 5.3 shows a nonparametric estimator of the regression function. Note that the vertical scale is different on that plot.

[^4]:    ${ }^{2}$ Generalized cross-validation has certain invariance properties not shared by leave-one-out cross-validation. In practice, however, the two are usually similar.
    ${ }^{3}$ Actually, this is not quite the $C_{p}$ formula. Usually, the estimate (5.86) for $\sigma^{2}$ is used.

[^5]:    ${ }^{4}$ More details on splines can be found in Wahba (1990).

[^6]:    ${ }^{5}$ If you are familiar with ridge regression then you will recognize this as being similar to ridge regression.

[^7]:    ${ }^{6}$ A different approach to estimating bias is discussed in Section 6.4 of Ruppert et al. (2003). However, I am not aware of any theoretical results to justify the resulting confidence bands.

[^8]:    ${ }^{7}$ This means it is a random function such that the vector $\left(W\left(x_{1}\right), \ldots, W\left(x_{k}\right)\right)$ has a multivariate Normal distribution, for any finite set of points $x_{1}, \ldots, x_{k}$.

[^9]:    ${ }^{8}$ That is, $\mathbb{P}\left(A_{1} \cup \cdots \cup A_{m}\right) \leq \sum_{i} \mathbb{P}\left(A_{i}\right)$.

[^10]:    ${ }^{9}$ It might be appropriate to use a weighted fit since the variance of the Bernoulli is a function of the mean.

[^11]:    ${ }^{1}$ Recall that the interquartile range is the 75 th percentile minus the 25 th percentile. The reason for dividing by 1.34 is that $Q / 1.34$ is a consistent estimate of $\sigma$ if the data are from a $N\left(\mu, \sigma^{2}\right)$.

[^12]:    ${ }^{1}$ The equality sign in (7.5) means that $\int_{a}^{b}\left(f(x)-f_{N}(x)\right)^{2} d x \rightarrow 0$ as $N \rightarrow \infty$, where $f_{N}=\sum_{j=1}^{N} \theta_{j} \phi_{j}(x)$.

[^13]:    ${ }^{2}$ The weak derivative is defined in the appendix.

[^14]:    ${ }^{3}$ Weak differentiability is defined in the appendix.

[^15]:    ${ }^{4}$ This is a finite-dimensional version of Pinsker's theorem. Theorem 7.32 is the usual version.

[^16]:    ${ }^{5}$ The Bayes estimator minimizes Bayes risk $\int R(\theta, \widehat{\theta}) d \pi(\theta)$ for a given prior $\pi$.

[^17]:    ${ }^{6}$ Intuitively, think of $d W(t)$ as a vector of Normals on a very fine grid.

[^18]:    ${ }^{1}$ More generally we could take $r_{n}(x)=\sum_{j=1}^{p(n)} \theta_{j} \phi_{j}(x)$ where $p(n) \rightarrow \infty$ at an appropriate rate.
    ${ }^{2}$ See Definition 7.2.

[^19]:    ${ }^{3}$ We call this a modified risk estimator since we have inserted an estimate $\widehat{\sigma}$ of $\sigma$ and we replaced $\left(Z_{j}^{2}-\widehat{\sigma}^{2} / n\right)$ with $\left(Z_{j}^{2}-\widehat{\sigma}^{2} / n\right)_{+}$which usually improves the risk estimate.

[^20]:    ${ }^{1}$ There are no smooth, symmetric, compactly supported wavelets.

[^21]:    ${ }^{2}$ In practice, we do not compute $S_{k}$ and $D_{j k}$ using (9.27). Instead, we proceed as follows. The highest level scaling coefficients $\alpha_{J-1, k}$ are approximated by $Y_{k}$. This is reasonable since $\phi_{J-1, k}$ is highly localized and therefore

    $$
    \mathbb{E}\left(Y_{k}\right)=f(k / n) \approx \int f(x) \phi_{J-1, k}(x) d x=\alpha_{J-1, k}
    $$

    Then we apply the cascade algorithm to get the rest of the coefficients; see the appendix for details. Some authors define

    $$
    S_{k}=\frac{1}{\sqrt{n}} \sum_{i} \phi_{j_{0}, k}\left(x_{i}\right) Y_{i} \quad \text { and } \quad D_{j k}=\frac{1}{\sqrt{n}} \sum_{i} \psi_{j k}\left(x_{i}\right) Y_{i}
    $$

    instead of using (9.27). This means that $S_{k} \approx N\left(\sqrt{n} \alpha_{j_{0}, k}, \sigma^{2}\right)$ and $D_{j k} \approx N\left(\sqrt{n} \beta_{j k}, \sigma^{2}\right)$. Hence, the final estimates should be divided by $\sqrt{n}$. Also, the estimate of the variance should be changed to

    $$
    \widehat{\sigma}=\frac{\operatorname{median}\left(\left|D_{J-1, k}-\operatorname{median}\left(D_{J-1, k}\right)\right|: k=0, \ldots, 2^{J-1}-1\right)}{0.6745}
    $$

[^22]:    ${ }^{3}$ In practice, SureShrink is sometimes modified to include one more step. If the coefficients at level $j$ are sparse, then universal thresholding is used instead. See Donoho and Johnstone (1995).

[^23]:    ${ }^{4}$ I've stated the result in a very specific form. The original result is more general than this.

[^24]:    ${ }^{1}$ See Chapter 11 of Wasserman (2004) for a general discussion of the advantages and pitfalls of Bayesian inference.

[^25]:    ${ }^{2}$ This is a poor choice of terminology. The Bayes rule has nothing to do with Bayesian statistical inference. Indeed, the Bayes rule $h$ can be estimated by frequentist or Bayesian methods.

